

NOTABLE COMPUTER NETWORKS

Computer networks are becoming more numerous and more diverse. Collectively, they constitute a worldwide metanetwork.

JOHN S. QUARTERMAN and JOSIAH C. HOSKINS

A *computer network* is a set of computers using common *protocols* to communicate over connecting transmission media. To warrant inclusion in this article, a network must provide at least mail or news service to its users and interconnect to other networks that provide such services.

Among these networks are *long-haul* (or *wide-area*) *networks* that can encompass continents. There are also *internets* of smaller networks communicating with one another through the same protocols. The internets of interest here include long-haul networks as constituents, and most also include *local-area networks* (LANs). The computers connected can be small microcomputers, supercomputers, or anything between. Multiple interconnection media can be used, including coaxial cable, optical fiber, satellite links, twisted pair, or telephone lines. The protocols can vary widely in speed, reliability, and general functionality. The services provided can range from the most basic mail service to distributed file systems and remote procedure call capability. A similar diversity applies to ownership, funding, administration, addressing, and other characteristics. Together, these networks form a *metanetwork* (sometimes called Worldnet) that is used daily by many communities of interest throughout the world.

This article is primarily concerned with describing specific networks. Various characteristics are given for each of them, where information was available. We also discuss some topics that are broader than

any one network, such as legal and social issues. In our taxonomy, there are five basic kinds of networks:

CHARACTERISTICS OF COMPUTER NETWORKS

Purpose, Administration, and Funding

In our taxonomy, there are five basic kinds of networks:

1. *Research Networks.* Many of the earliest computer networks were designed and implemented for research in computer networking. There are still a number of networks that are either themselves research projects or are administered in support of other research. The ARPANET is the best-known example. Such networks are usually administered by government agencies or contractors, and supported by government grants. Their users and host machines do not ordinarily pay directly for service. Because of their goals and the nature of their funding, access to these networks tends to be limited to researchers participating in the funded work.

There are also a number of military networks, mostly either of the same technology or closely interoperable with certain research networks. In fact, many of the research networks had as one of their goals the eventual development of corresponding military networks. The best-known example of this sort of pairing is ARPANET and MILNET, which are closely allied in the ARPA Internet. Because of the

© 1986 ACM 0001-0782/86/1000-0932 75c

close associations (and because what can be written about them is too limited to fill a separate category), military networks are included in the research network category.

There are also networks that were developed to provide ARPANET-like services to people and organizations who could not obtain access to the government-sponsored military or research networks, or who did not wish to be associated with military work. The best-known of these (among those of interest in this article) is CSNET. Users of those networks usually pay an annual connection fee.

2. *Company Networks.* Large corporations like Xerox, DEC, IBM, and AT&T have implemented internal networks in support of their business operations. Many of these are just LAN within particular buildings, although some are international or even intercontinental in scope. The administration and

funding of such networks usually come from a single company, and their users are mostly employees of that company.

3. *Cooperative Networks.* These are networks that have grown up among communities of users with similar interests. Many of them, such as BITNET and its associated networks NETNORTH and EARN, originated in an academic environment. Some originated among users of a particular vendor's systems (e.g., BITNET and IBM) or a particular operating system (UUCP, USENET, EUnet, JUNET, and ACSNET among UNIX[®] users) or both (FidoNet among IBM PC and MS/DOS users). Many, such as ACSNET, EUnet, JUNET, UUCP, and USENET, have users with a mixture of academic, corporate, research, and commercial interests. Often the strongest bond in a particular cooperative network (at least

UNIX is a trademark of AT&T Bell Laboratories.

Contents

CHARACTERISTICS OF COMPUTER NETWORKS.....	932	ROSE	950
Purpose, Administration, and Funding	932	COMPANY NETWORKS.....	951
Layers, Protocols, and Services	934	Xerox Internet.....	951
Layering Models	934	DEC's Easynet	951
Application Protocols	935	IBM's VNET	952
Presentation and Session Protocols	937	AT&T	953
Transport Protocols	937	COOPERATIVE NETWORKS.....	953
Networks	937	BITNET	953
Reliability	937	BITNET in the United States	953
Addressing and Routing	938	NETNORTH	955
Security	939	EARN	955
Costs	939	FidoNet	955
Scope	941	ACSNET	955
UUCP	941	UUCP	957
USENET	941	USENET	958
EUnet	942	EUnet	959
JUNET	942	SDN	961
TUNET	942	JUNET	962
METANETWORKS	942	METANETWORKS	963
CSNET	943	CSNET	963
NSNET	944	NSNET	963
National Research Internet	944	National Research Internet	964
KARE	944	KARE	964
Pacnet	945	Pacnet	964
ADSEANet	945	ADSEANet	964
BULLETIN BOARDS AND NETWORKS	945	BULLETIN BOARDS AND NETWORKS	964
Commercial Networks	945	Commercial Networks	965
Conferencing Networks	945	Conferencing Networks	965
SOCIAL AND LEGAL ISSUES	946	SOCIAL AND LEGAL ISSUES	966
Networked Communities	947	Networked Communities	966
Technical Groups	947	Technical Groups	966
Social Groups	948	Social Groups	967
Social Effects	948	Social Effects	967
Legal Issues	949	Legal Issues	967
HISTORY	949	HISTORY	967
BIBLIOGRAPHIC NOTES	950	BIBLIOGRAPHIC NOTES	968

Notable Acronyms from "Notable Computer Networks"

ARPA	Advanced Research Projects Agency—the acronym has been changed to DARPA, for Defense Advanced Research Projects Agency
CCITT	International Consultative Committee on Telegraph and Telephony (from the French)
DDN	Defense Data Network—the DDN PMO is the Defense Data Network Program Management Office
FTP	File Transfer Protocol
FTAM	File Transfer and Management
IMP	Interface Message Processor
IP	Internet Protocol
ISO	International Organization for Standardization
NCP	Network Control Protocol
PAD	Packet assembler/disassembler
PSN	Packet-Switch Node
RFC	Request for Comments—the RFCs are a set of on-line documents providing information about the ARPA Internet
RSCS	Remote Spooling and Communications Subsystem—this is the spooling protocol used in VNET and BITNET
SMTP	Simple Mail-Transfer Protocol
TCP	Transmission Control Protocol
TP0-TP4	The set of transmission protocols in the ISO protocol suite
UDP	Uniform Datagram Protocol
UIP	User Interface Presentation
UUCP	UNIX to UNIX CoPy
X.25	The network layer protocol in the ISO protocol suite
X.400	The ISO mail protocol
XNS	Xerox Network System

initially) is the network or transport protocol used (e.g., RSCS for BITNET, UUCP for the UUCP mail network, and FIDO for FidoNet).

Administration of cooperative networks is generally distributed. Some, like BITNET, have a certain amount of centralized control and organization. Others, like UUCP, function in near-complete anarchy. Fees are not generally collected by a central organization, but are paid by each node for connections to other nodes (for example, in telephone bills). Some cooperative networks nonetheless receive strong infusions of money from specific companies. For instance, BITNET has received massive contributions from IBM, AT&T spends more on UUCP and USENET than any other company (though there are others that also pay far more than their share), and DEC and Philips have provided strong backing for EUnet.

4. *Commercial Networks.* These networks provide services to outside users for profit. Some are well-known public data networks, such as TYMNET and TELENET. Many are common carriers like the telephone system. Administration is always centralized, though execution may be delegated. Funding is usually derived from fees charged to individual persons or organizations for connect time or CPU time. Most commercial networks do not permit free exchange of mail with other networks and are thus outside the scope of this article.

5. *Metanetworks.* There are several projects being directed at extending network user communities by connecting existing and as yet unconstructed networks into metanetworks. Such networks differ from internetworks both in their goals and in that their constituent parts often have dissimilar protocols even as high as the transport layer. Most metanetworks do not yet exist: CSNET is an exception.

Several networks do not fall neatly into any of these categories: COSAC, JUNET, and SDN could be classified as either research or cooperative. COSAC is listed as a research network because it is related to the EAN networks and DFN. SDN and JUNET are grouped with the cooperative networks because they are related to UUCP, USENET, and EUnet.

Layers, Protocols, and Services

Computer network protocols can be quite complex. To keep complexity manageable, protocols are designed in layers, building up from those near the hardware to those near the users. In each layer, there may be one or more protocols that peer entities on that layer can use to communicate with one another. The interfaces between adjacent layers are defined, and protocol designers often assume that nonadjacent layers do not communicate directly [18].

Layering Models. The International Organization for Standardization (ISO) has proposed a standard reference model for what they call Open Systems Interconnection [39]. This model has seven basic layers: physical, data link, network, transport, session, presentation, and application. The network layer is often assumed to be X.25, a protocol in a series promulgated by CCITT. The transport protocols, TP0 through TP4, provide different classes of service, ranging from simple datagrams to reliable connections (TP2 is designed especially for use over X.25). The higher layers are nearing design completion, and many of them are already implemented.

Much of the ISO work is based on the work of those who designed and continue to do research on

the ARPANET and the ARPA Internet [8], as well as on related early network efforts such as CYCLADES [72]. The ARPANET originally had three basic layers: network, transport, and process/applications, plus the network hardware. The ARPA Internet adds a fourth, internet layer, for which the IP (internet) protocol is used. (There is also a physical layer, and some descriptions distinguish a link layer plus a utility layer that is similar to a combination of the ISO presentation and session layers.) ISO has also recently adopted an internet sublayer of the network layer that strongly resembles IP. The two most commonly used transport protocols in the ARPA Internet are transmission control protocols (TCPs) [53] (reliable connections) and uniform datagram protocols (UDPs) [83] (datagrams). There is contention between the proponents of the ISO Reference Model (ISORM) and the ARPANET Reference Model (ARM) or ARPA Internet Reference Model [8]. The ISO protocols expect virtual circuits at the network layer, whereas the TCP/IP suite makes more use of datagrams (IP is a datagram protocol). (This is an exaggeration for pedagogical purposes of the true situation.) Table I shows some of the differences in layering as well as the layers used in a networking implementation in the 4.3BSD version [73] of the UNIX operating system.

There is a third major protocol suite called Xerox Network Services (Xerox NS) (see the "Xerox" section), having layers that are similar to the above models. There are also the Read Book protocols, which are primarily used in the United Kingdom in JANET. For the purposes of this article, the four layers—network, internet, transport, and applications—are at about the right level of detail.

The protocol suites discussed above were designed with the assumption of dedicated links between network nodes. There are other protocols, such as UUCP and Fido, that were designed for use with intermittent connections. These will be discussed below in sections on the networks that use them, along with other protocols for dedicated connections.

Application Protocols. The types of services (application protocols) provided by the various networks vary widely, but they tend to fall into recognizable classes.

Electronic mail is the most widespread and the most rudimentary. It allows a user to send a message to another user on either the same or a different host. The message is placed in the recipient's mailbox on the destination host. Mail is typically a one-to-one point-to-point communication medium, though it is possible to mail to more than one recipient at the same time. Most networks implement mail in a batched, asynchronous manner, with errors being reported by mail messages from a mail daemon process. Ordinarily, anyone can send mail to anyone with computer access.

File transfer (sometimes known as FTP or FTAM) is probably the second most common service. It allows files to be transferred from one host machine to another. Since data formats vary widely among operating systems and machine types, there are usually several file transfer formats supported. The most generally usable one is ASCII text. The user initiating a transfer must have read access on the source file and write or create access on the destination file or directory.

Remote command execution (a special case of which is Remote Job Execution or RJE) facilities are pro-

TABLE I. Network Reference Models and Layering

Protocol layers from the ISO and ARPANET reference models with examples of layering implementations in the 4.3BSD				
ISO Reference Model	ARPA Internet layers	4.3BSD implementation layers	Examples of uses of layer in 4.3BSD	
Application	Process/Applications	User programs and libraries	telnet, ftp, rlogin, or rcp	named, time, rwho, or talk
Presentation				
Session		Sockets	SOCK_STREAM	SOCK_DGRAM
Transport	Transport		TCP	UDP
Network (Internet)	Internet	Protocols	IP	
Data link	Network	Network interfaces	Ethernet driver	
Physical	Physical	Network hardware	Interlan controller	

vided about as frequently as file transfer. They allow a user to execute a command on a remote machine. There are usually either strong permission checks or a sharply limited set of commands that can be executed.

Remote login is the simplest service in concept, though many networks fail to provide it. It allows a user to access a system on a remote host over a network as if connected to a direct terminal. There

are some complications because different systems expect different terminal types, and some networks distinguish remote terminal access from host-to-host remote logins. Permission checks are usually similar to those for access by a direct terminal line.

Computer conferencing systems are message exchange systems that are generally similar to electronic mail systems, but differing in that they are many-to-many (or broadcast) rather than one-to-one

TABLE II. Characteristics of Some Notable Networks (see Table Ila for an explanation of the symbols used here)

Names	Center	Extent	Hosts	Users	Layers	Services	Quality
<i>Research networks</i>							
ARPA Internet	USA	3.8	2,050	?	i T	lfmo	56,000
ARPANET	USA	1.1	150	?	a T	lfmo	56,000
MILNET	USA	2.3	400	?	a T	lfmo	56,000
MINET	Europe	1.4	?	?	x T	lfmo	9600
CSNET	USA	4.10	170	?	i C	m	1200
Phonenet	USA	4	128	?	d M	m	1200
X25NET	USA	1	18	?	x T	lfm	9600
ARPANET (CSNET hosts)	USA	1.1	25	?	i T	lfmo	56,000
Cypress	USA	1	6	?	p T	lfmo	9600
NSFNET	USA	1.1	65	?	i T	lfmo	56,000
MFENET	USA	1.1	120	?	i D	fm	56,000
SPAN	USA	1.1	100+	?	px D	fm	56,000
MAILNET	USA	2.3?	28	1,800	dx M	m	1200
JANET	UK	1.1	915	?	x B	lfmo	4800
<i>European networks</i>							
CDNnet	Europe	3.12	33	?	x X	m	2400
COSAC	Canada	1.1	32	?	x X	lfmo	2400
DFN	France	1.2	27	?	x X	fm	1200
	Germany	1.1	6	?	x X	fm	9600
<i>Company networks</i>							
Xerox Internet	USA	3.4	?	12,000	i N	lfmo	56,000
Xerox RIN	USA	1	?	4,000	i NPT	lfmo	56,000
Xerox CIN	USA	3	?	8,000	i N	lfmo	56,000
DEC's Easynet	USA	4.20?	10,000+	60,000?	d D	lfmo	56,000
IBM's VNET	USA	4.?	2,200	?	p R	fmo	9600
<i>Cooperative networks</i>							
BITNET	USA	3.21	1,306	?	p R	fmo	9600
BITNET	USA	1.2	845	?	p R	fmo	9600
NETNORTH	Canada	1.1	91	?	p R	fmo	9600
EARN	Europe	1.17	363	?	px R	fmo	9600
Asianet	Japan	1.1	7	?	p R	fmo	9600
FidoNet	USA	2.4?	500	?	d F	m	1200
ACSNET	Australia	1.1	300	?	depx A	fmno	1200
UUCP mail	North America	4.5	7,000+	200,000?	dx U	m	1200
USENET news	North America	4.4	2,500+	50,000?	dx UATR	n	1200
EUnet	Europe	1.13	896	?	dpx UAT	lfmo	1200
SDN	Korea	1.1	100	?	epx TU	lfmo	2400
JUNET	Japan	1.1	160	?	x U	mn	2400

TABLE IIa. Legend for Table II.

Extent:	A pair of the numbers of continents and nations reached.
Layers:	Network layer (left column) and internet or transport layers (right column).
Protocols for the network layer:	
a:	ARPANET-style PSN (BBN 1822) communications subnet.
d:	Dial-up telephone line.
e:	An ethernet.
i:	An internet over various network layers.
p:	Leased telephone line.
x:	X.25 (usually over leased telephone line).
Protocols for the internet and transport layers:	
A:	ACSNET's SUN-II.
B:	JANET's Coloured Books.
C:	CSNET's TCP/IP (ARPANET), TCP/IP on X.25 (X25NET), MMDf2, (Phononet).
D:	DEC's DECNET.
F:	FidoNet.
N:	Xerox Network Services protocol suite.
P:	Xerox PARC Universal Packet (PUP) protocol.
T:	ARPA's TCP/IP protocol suite.
U:	AT&T's UUCP.
R:	IBM's RSCS.
X:	CCITT/ISO X.400 and related protocols.
Services:	l: remote login; f: file transfer; m: mail; n: news; o: other.
Quality:	A triple of speed, delivery, and reliability.
Speed:	Most typical speed of long-haul links in bits per second (bps): 300, 1,200, 2,400, 9,600, 19,200, 56,000, 1.1 microwave.
Delivery:	Average delivery time for mail messages: m: minutes; h: an hour or more; d: a day or more.
Reliability:	On a subjective scale from 1 (lowest) to 10 (highest).

la. Usually one copy of a message is kept per rather than one per user as for mail. Though conferencing systems are synchronous and immediately interactive, most of the ones mentioned in article are batched and asynchronous, since they are distributed over wide areas. USENET news bulletin-board systems are examples of this kind of system. (See the "Bulletin Boards and Networks" section.)

Some networks support sophisticated services such as remote procedure call, distributed databases, network file systems. Most of these work best on local networks, however, and most long-haul networks are not fast enough. Such services are not fully supported on the networks we discuss.

Presentation and Session Protocols. The ISO model distinguishes a presentation layer and a session layer. These are of more concern to programmers than to users of the networks, and many of the networks we discuss do not distinguish such layers clearly. We thus have little to say about them.

Transport and Internet Protocols. There are a number of protocols that are used either on more than one network or on large, significant networks. We provide a few details on such protocols in the description

of one of the main networks on which they are used.

Network Protocols. There is not sufficient space to discuss network layer protocols at any length. Tables IIa and IIb indicate which ones are used by which networks.

Hardware. The main concern with hardware from our perspective is the speed and reliability it can confer on any particular network.

Speed and Reliability

It is difficult to find metrics of speed and reliability that can be applied to a range of networks as diverse as that covered in this article. It is even more difficult to get enough information to apply such metrics. We provide a few ad hoc measures and do not pretend to treat this topic in depth. For reliability, we use a subjective scale from 1 (worst) to 10 (best). For speed, we use two measures: bits per second (bps) and average time of delivery of mail.

Though many internets may include ethernets, ring networks, or other fast local-area networks, we use bps for the most common long-haul links between widely separated hosts. By "most common" we mean most likely to be used in ordinary commu-

nications, not most commonly implemented. If a network had many 9,600-bps links and one T1 microwave link, and most traffic used the latter, we would list the T1 speed.

The average time for delivery of mail varies so much even within networks that we list only three values:

- *minutes*, indicating a delivery time of less than an hour;
- *hours*, indicating a delivery time of at least an hour, but less than a day;
- *days*, indicating a delivery time of at least a day.

We also record whether a network consists mostly of dial-up telephone links or of dedicated links.

Naming, Addressing, and Routing

These three related terms, which are important to networking, are often confused [46, 77]. The *name* of a host, mailbox, or other resource is what a user uses to indicate the resource desired. Its *address* specifies the location of the resource to the network software. A *route* is used by the network software to determine how to get to the resource. In the public switched telephone network, a name is a personal name, such as Jane Doe, an address is a telephone number, and a route is a sequence of telephone lines and exchanges that are used to reach Jane's number from the caller's telephone.

Consider hosts on the ARPA Internet. A host might be named SALLY.UTEXAS.EDU and have an Internet address of 10.2.0.62. The address would be discovered by the software on the user's machine (either by old-style static host table lookup or by new-style domain nameserver protocols). The IP protocol would then use the address to route the packet to the appropriate network. The network named by the address 10.2.0.62 is network 10, the ARPANET. The ARPANET has a communications subnet of computers called Packet-Switch Nodes (PSNs) to which hosts are attached. The PSNs then extract an ARPANET address (host 2 on PSN 62) from the IP address and use it to determine a route to the destination PSN and thus to the destination host. Note that names and addresses are relative to network protocols. The IP address is treated as a name when the ARPANET address is extracted from it. Routing is first done on the IP address and then on the ARPANET address.

Naming, addressing, and routing can all be hierarchical. SALLY.UTEXAS.EDU is an ARPA Internet domain name, where EDU is a top-level domain, UTEXAS.EDU a subdomain of EDU, and

SALLY.UTEXAS.EDU a further subdomain (in this case, SALLY.UTEXAS.EDU is a host machine). The user interface software on machines in the UTEXAS.EDU domain may allow users to abbreviate SALLY.UTEXAS.EDU as SALLY. However, there could be another host named SALLY.CSS.GOV, in which case the abbreviation SALLY on hosts in the domain CSS.GOV would not refer to the same host as in UTEXAS.EDU.

The address 10.2.0.62 is actually a two-level ARPA Internet address. The prefix 10 is the network number of the ARPANET, and the rest (the local part) is a host number on the ARPANET. The local part can be mapped to a network address by different methods for different networks. In this particular case the network address is actually contained in the Internet address, and there is a further hierarchy to the host address. The final 62 is the PSN number, and the rest is the host-on-PSN number.

Routing in the Internet is also hierarchical. If a route is found to the appropriate network through gateways (by the Gateway to Gateway Protocol (GGP) [33] and the Exterior Gateway Protocol [54, 76]), then a route is found to the appropriate host on the network (by protocols appropriate to the network). In the ARPANET, the latter problem reduces to finding the host's PSN, the number of which is encoded in the address. For an address on an ethernet (e.g., 128.83.138.11), finding the appropriate host is usually simpler since ethernet is a broadcast network. (Ethernet (capital E) is a Xerox protocol used for LAN, whereas an ethernet (small e) refers to an Ethernet-like network.)

There are two kinds of routing: *source routing*, where the user supplies the route to the desired resource, and *system routing*, where the network software determines a route. Most networks and internets provide system routing [26]. There are a few exceptions, most prominently UUCP. The metanetwork of differing networks and internets frequently requires source routing to reach the appropriate network because there is as yet no universally accepted network addressing convention. Source routes like "alpha!beta%gamma@delta" are thus unfortunately still common.

A resource may have more than one name, address, or route. In the ARPA Internet, SALLY.UTEXAS.EDU might have two addresses, 10.2.0.62 and 128.83.138.11, if it were connected to two networks. Though hosts in the Internet have only one primary name, they may be known by other names on non-Internet networks. For instance, SALLY.UTEXAS.EDU might be known as ut-sally on the UUCP network. It would be better if every host

had one name for all networks, but that is not yet possible. Both the IP protocol and the ARPANET network protocol are datagram based, and different datagrams can pass through different routes to reach the same destination, even when the source is the same.

Domains. The ARPA Internet domain name system is an attempt to decentralize administration of the mapping of host names to host addresses by the use of nameservers, each of which controls part of the name space [57, 59, 70, 71]. This became necessary partly because the static host table formerly used for that purpose had become unwieldy with the growth of the Internet and partly because most of the hosts in the Internet are on networks local to particular organizations thus making it desirable to allow the local administration to control that mapping. The domain name system also implements a hierarchical naming scheme and provides protocols for communication with the nameservers [55, 56, 58, 67].

The British network JANET has a domain system similar to that of the ARPA Internet, but with the domains in the opposite order. The root is on the left rather than on the right. The Australian network CSNET also has an Internet-like domain name system.

At a recent meeting, North American representatives of the ARPA Internet, BITNET, CSNET, and UUCP decided to adopt the ARPA Internet domain naming syntax and domains as a common naming scheme [68, 69]. (The adoption is voluntary on a per-host basis on UUCP and BITNET.) EUnet in Europe is going in the same direction and has already defined several top-level national domains. JUNET in Japan already has a similar domain system. A network, NSFnet, has also standardized on UUCP and related protocols. Thus Internet domain names may become the de facto standard, at least in the United States.

ISO X.400 mail standard also has a domain system which uses *attributes*. A resource is defined by name and several attributes. Name conflicts can be resolved by specifying sufficient attributes [39]. There is a similar mechanism at the network level [175].

The EAN networks use a simplified version of the X.400 system because there is as yet no registry for X.400 domains.

Source Routing. Possibly the only widespread network without system routing is UUCP. Practically everyone agrees that users should not have to supply source routes manually. The usual means of avoiding this on the UUCP network is for the system

administrator to run a program (*pathalias*) on a large database of information about most of the hosts in the network and which other hosts they connect to. The result is a database of source routes from the local machine to all the machines that are in the connectivity database. The local mail system uses the source route database to convert addresses into routes.

Problems with this scheme include the extravagant use of CPU time and disk space required to run *pathalias* and keep copies of the databases, plus the fact that the connectivity database is published only once a month and is thus guaranteed to always be out-of-date.

Attribute Lists. Names and addresses can be either absolute or relative. In the ARPA Internet, both Internet addresses and fully qualified domain names are absolute (within the Internet), but user mailbox names are relative to domain names. Most other networks have absolute names and addresses (again, UUCP is an exception).

Relative names are a problem because they make mapping into addresses ambiguous. This is why short names like SALLY are considered to be only abbreviations for a single primary name such as SALLY.UTEXAS.EDU; it is the responsibility of the local user interface to produce the primary name when communicating with any other host. Relative addresses are a problem because a host may have a different address depending on where it is being addressed from. Both relative names or relative addresses leave open the possibility that two hosts might have the same address, which would make proper routing impossible. Nonetheless, maintaining absolute names is difficult, since absolute really means relative to some standard, and there is no universal standard. X.400 is one attempt to handle this problem.

The UUCP network has not had absolute host names or addresses. A single name (e.g., bilbo) may be assigned by several different companies to several different machines. This may happen because a company was not connected to the general UUCP network at the time and thus was unaware of the conflict, or because a host was not originally expected to communicate with the world at large, or because the first bilbo was not listed in the UUCP map, or for other reasons.

One method for disambiguating such conflicts is to refer to each bilbo by a route from a well-known neighbor (e.g., princeton!bilbo or ihnp4!bilbo). These partial routes are a kind of attribute list in the X.400 sense. Of course, if someone names another host

TABLE III. Address Formats and Gateways

Domains:		com, edu, etc.		CSNET Phonenet		MAILNET	
From	To	From	To	From	To	From	To
ARPA Internet	u@d.i	u%h.csnet@relay.cs.net	u%h.mailnet@mit-multics				
CSNET Phonenet	u@d.i	u@d.i	u%h.mailnet@mit-multics				
MAILNET	u%d.i@mit-multics	u%d.i%relay.cs.net@mit-multics	?				
JANET	u%d.i@uk.ac.ucl.cs	?	?				
EAN	u@d.i	u@d.i	u@h.mailnet				
COSAC	adi/u%d.i@relay.cs.net	adi/u%d.i@relay.cs.net	?				
BITNET	u@d.i	u@h.csnet	u@h.mailnet				
ACSNET	u%d.i@munnari.oz	u%d.i@munnari.oz	u%h.mailnet@munnari.oz				
UUCP	g!d.!u	g!d.!u	g!h.mailnet!u				
JUNET	u@d.i.arpa	u@d.i.arpa	u%h.mailnet@mit-multics				
Domains:		U. uk		com, edu, etc.		COSAC	
From	To	From	To	From	To	From	To
ARPA Internet	u%d.U@cs.ucl.ac.uk	u%d.E@ubc.csnet@relay.cs.net	h/u%france.csnet@rela				
CSNET Phonenet	u%d.U@cs.ucl.ac.uk	u%d.E@ubc.csnet	h/u@france.csnet				
MAILNET	?	u%d.E@ubc.mailnet	?				
JANET	u@U.d	u@d.U	?				
EAN	u@d.U	u@d.E	h/u@france.csnet				
COSAC	?	?	h/u				
BITNET	u%d.U@ac.uk	u@d.E	h/u@france.csnet				
ACSNET	u%d.U@munnari.oz	u%d.E@munnari.oz	h/u%france.csnet@rela				
UUCP	g!cs.ucl.ac.uk!d.U!u	g!d.E!u	?				
JUNET	u@d.U.janet	u%d.E@ubc.csnet	h/u@france.csnet				
Domains:		R. A registry		DEC's EasyNet		IBM's vnet	
From	To	From	To	From	To	From	To
ARPA Internet	u.R@xerox.com	u%h.dec@decwrl.dec.com	u%h@ibm.com				
CSNET Phonenet	u.R@xerox.com	u%h.dec@decwrl.dec.com	u%h@ibm.com				
MAILNET	?	?	?				
JANET	?	?	?				
EAN	u.R@xerox.com	u%h.dec@decwrl.dec.com	u%h@ibm.com				
COSAC	?	?	?				
BITNET	u.R@xerox.com	u%h.dec.com@decwrl.dec.com	u@vnet				
ACSNET	u.R%xerox.com@munnari.oz	u%h.dec.com@munnari.oz	u%h%ibm.com@munnari.oz				
UUCP	parcvax!u.R	decwrl!h.dec.com!u	g!ibm.com!u%h				
JUNET	u.R@xerox.com.arpa	u%h.dec@decwrl.dec.com.arpa	u%h@ibm.com.arpa				
Domains:		At. oz.au		JUNET		JUNET	
From	To	From	To	From	To	From	To
ARPA Internet	u%h.bitnet@wiscvm.wisc.edu	u@d.A	u%h.uucp@g	u%d.J%utokyo-relay@			
CSNET Phonenet	u%h.bitnet@relay.cs.net	u@d.A	u%h.uucp@g	u%d.J@utokyo-relay			
MAILNET	?	u%d.A%g@mit-multics	?	u%d.J%csnet-relay@mit			
JANET	?	u%d.oz@uk.ac.ukc	?	u%d.J@uk.ac.ukc			
EAN	u@h.bitnet	u@d.A	u@h.uucp	u%d.J@relay.cs.net			
COSAC	adi/u%h.bitnet@relay.cs.net	?	adi/u%h.uucp	adi/u%h.J@relay.cs.net			
BITNET	u@h	u%d.A@g	h1!h2!h!u@psuvax1	u%d.J@csnet-relay.csnet			
ACSNET	u%h.bitnet@munnari.oz	u@d.A	u%h.uucp@munnari.oz	u%d.J@munnari.oz			
UUCP	psuvax1!h.bitnet!u	seismo!munnari!d.A!u	h1!h2!h!u	g!d.J!u			
JUNET	u@h.bitnet	u@d.A	u@h.uucp	u@d.J			

Notes: From UUCP to CDNnet *ubc-can* is a gateway; from EUnet to the European EAN networks there is one gateway per country; there is more than one gateway between BITNET and UUCP; UUCP, EUnet, and SON are similarly addressed, so only one of them is listed here.
Abbreviations: *u*: user; *h*: host; *g*: gateway (unnamed here); *d*: domain.
Omissions: From commercial networks, with commercial networks, with the ARPA Internet.

"princeton," or if princeton leaves the network, a longer or different partial route would have to be given for that bilbo. This problem occurs with all attribute list schemes: Names and addresses are not absolute.

Another possible solution, now being worked on by a group called the UUCP Project, is to give each UUCP host an ARPA Internet domain name, such as bilbo.princeton.edu. The former UUCP name would still be used as a kind of network address. Routing would be done from domain to domain, so network-wide tables would only be needed for routes to domain gateway hosts, and complete connectivity information would only be kept on hosts within a subdomain by those same hosts (similar methods are already used in EUnet). The UUCP network would thus be integrated into the ARPA Internet domain name system. This plan is opposed by some people who actually like UUCP source routing. For an interesting discussion of related issues by a prominent person on each side, see [1]. We should point out that source routing, attribute lists, and domain names are mutually exclusive—at least not on the UUCP network. Each can be used in combination with the

others. There are several related and somewhat controversial terms related to machines that interconnect networks. These include repeaters, bridges, routers, relays, and gateways. Most of these are below the upper layers of protocols and are transparent to the users. Here we are concerned with gateways between networks with dissimilar internal layers. These usually work less well than relays at lower layers, are often less transparent,

and usually have to be considered by the user when sending mail across such network boundaries. (Mail is often the only service that can be used.) In some cases, such gateways may not be known. In others, it may not be possible to reveal them because of political or economic considerations. Table III is a compilation of likely mail routing syntaxes between some of the networks discussed in this article. However, gateways are subject to change, and the nature of the information makes it impossible to compile a table that will be accurate for very long. In particular, addresses using a percent sign (%) to indicate indirection through a relay host (a kind of source routing) are a kludge that most people hope will be temporary. A specification (RFC987 [45]) has recently been formulated for translation between ARPA Internet domain addresses and X.400 attribute addresses. Software now exists to do that translation and also to translate between X.400 and EAN addresses. When such software is in general use, percent sign source routing should no longer be necessary between those kinds of networks.

The user interface may even vary among systems on the same network. The examples in Figure 1 (courtesy of Christian Huitema and Steve Kille) address the same person.

User interface presentation (UIP) refers to the representation of an address to the user. The first three examples are for networks whose internal naming formats use ASCII text and also are the same as the UIP. The next five examples represent the same binary X.400 encoding, and the last two represent the same EAN address. The binary encoding of X.400 addresses allows all networks that use it to communicate, but there is no single standard human-

```

@cs.ucl.ac.uk
@uksteve
@cs.ucl.ac.uk
des/steve(ucl/cs)
GB/ADMD=BT/PRMD=DES/O=UCL/
=CS/S=Kille/
b.A=bt.P=des;O=ucl;S=steve;OU=cs)
uclcs&des%bt&gb
uclcs#des&bt.gb
@cs.ucl.des.bt.gb
=ADMD=/PRMD=UK/DD=cs.ucl.ac/
DD=steve/

```

```

Via ARPA Internet
Via UUCP
Via JANET
X.400, GIPSI (of INRIA) UIP
X.400, RFC987 UIP
X.400, another UIP
X.400, DFN UIP
X.400, EARN/X.400 gateway UIP
EAN, RFC822 UIP and domain order
EAN, X.400 encoding, RFC987 UIP

```

FIGURE 1. Sample Addresses for Different Networks

readable text UIP. Confusion results from different user interface software, from differing addressing syntaxes peculiar to specific networks, from attempts to represent one network's syntax in another's, and from attempts to encapsulate one network's syntax inside another's.

The moral of all this is that there is no magic formula to get mail between any two points in Worldnet. It's a jungle with trails that may cross and conflict, lead to the wrong place, or become overgrown.

Size and Scope

It is difficult to find a single metric for size that is meaningful on all networks. The traditional unit is number of hosts. This is useful for networks like ARPANET or CSNET, where most nodes are medium-size time-sharing systems and the exact number of users on each is hard to determine. Some networks consist primarily of workstations (Xerox Internet) or personal computers (FidoNet) where there is usually one user per host (though many FidoNet nodes are bulletin boards that may have many users). Others, such as BITNET and its relatives, consist mostly of large IBM and Digital Equipment Corporation (DEC[®]) mainframes that are hosts in the ARPANET sense, but have many more users per host. Also, the number of users who have access to a network is not usually the same as the number who actually use the network. Thus the number of active mailboxes, for instance, may be interesting, but is usually hard to determine.

The most common unit of measurement we use in this article is number of hosts. Where possible, we also give number of users or such other measures as we can find.

It is useful to distinguish several common terms as used herein:

- A *machine* is a computer of any size.
- A *system* is a computer system of any size; this term is usually used synonymously with machine.
- A *node* is any vertex of a graph representing a network, that is, any machine (or system) on a network.
- A *host* is a network node that has resources of its own (such as disks, user mailboxes, or user accounts). A host is not a node (such as an X.25 PAD or an ARPANET TAC) used only to connect across the network to other nodes. Nor is a host a gateway. A single machine or system may serve both as a host and as a gateway, however.

DEC is a trademark of Digital Equipment Corporation.

- Finally, a *site* is a place (such as a building, company, or campus) where a group of network nodes is located. However, the term has a more specific meaning in CSNET, whereas UUCP, USENET, and EUnet users often use the word "site" to refer to a host. To avoid confusion, the latter usage does not occur in this article.

Most of the networks we describe have wide geographical extent, but the distribution of their hosts (or users) is not uniform. Many of the internets consist of many local-area networks connected by a few long-haul networks. Thus the hosts cluster on the local-area networks, which themselves tend to cluster. Most networks in North America have concentrations of hosts in Silicon Valley near San Francisco, Route 128 near Boston, and in the Toronto area because computing-related companies in North America tend to be concentrated in these areas. Primarily academic networks such as MAILNET and CSNET are widely dispersed, with nodes mostly at academic institutions. USENET and UUCP have concentrations in New Jersey because of AT&T.

Access

Networks generally have rules (or at least guidelines) controlling access to their services. In the descriptions of the networks, we list a network administrator to contact for information regarding access guidelines and more detailed information about a specific network whenever possible. For some networks, we have no published references to access; in many such cases, someone associated with the network's administrative organization was the source.

RESEARCH NETWORKS

ARPA Internet

The ARPA Internet is an internetwork of several networks all running the TCP/IP protocol suite [5] connected through gateways, and sharing common name and address spaces [8]. The ARPANET is the oldest of the networks in the ARPA Internet. Both are named after the Defense Advanced Research Projects Agency (DARPA), which is part of the U.S. Department of Defense. DARPA (formerly known as ARPA) has long been a major sponsor of networking research.

Internet with a capital *I* refers to a specific internetwork, usually the ARPA Internet, whereas internet with a small *i* can refer to any internetwork. (There is also the Xerox Internet, which uses different protocols and may be older.) Other networks, such as BITNET, UUCP, EUnet, and ACSNET, are not part of the capital-*I* Internet. CSNET is a special case: Part

of it is part of the ARPA Internet, and part of it is not.

The ARPA Internet exists to facilitate sharing of resources at participating organizations and collaboration among researchers, as well as to provide a testbed for new developments in networking. Practical coordination of the entire Internet is provided by the Network Information Center (NIC) at SRI International and the Network Operations Center (NOC) at Bolt, Beranek and Newman (BBN).

The two main backbone networks of the Internet, ARPANET and MILNET, are funded mostly by government grants. The campus area networks are mostly funded by local organizations. There are in general no per-user or per-message charges. Services include remote login (telnet), file transfer (FTP), mail (SMTP), and numerous other smaller services (gate, time, system status, Internet directory, etc.). ARPANET and MILNET hosts are all connected to a subnet of PSNs, which are then connected to each other over 56,000-bps dedicated lines, plus a few satellite links (e.g., to Hawaii). PSNs were formerly known as IMPs (for interface message processors). Reliability is usually very high, but may suffer during implementation of new capabilities (e.g., during implementation of the domain name system). It suffers during peak periods, and telnet can be painful then, but mail always gets through in a reasonable amount of time, usually in minutes. The old top-level domain (ARPA) is temporary and will vanish soon. Many hosts are already registered in new domains, which are

- commercial organizations;
- educational organizations;
- civilian government organizations;
- Department of Defense;
- administrative organizations for networks such as CSNET, UUCP, and BITNET;
- other organizations.

Some networks may be in the same domain (as at universities), and a single network may have several domains (as does the ARPANET). There are also domains for countries, such as UK for United Kingdom and AU for Australia. There are many people outside the United States (and some within) who claim that all of COM, EDU, GOV, etc., should be under a top-level domain (U.S.). ARPANET and MILNET have about 150 and 400 hosts, respectively. There are also numerous networks at universities and private companies that are part of the Internet. Since many of the hosts on such local networks are known only to their parent organizations, the real size of the Internet is hard to judge.

However, there are probably more than 2000 hosts and tens to hundreds of thousands of users. There are special nodes called TACs (Terminal Access Controllers) on both ARPANET and MILNET whose only function is to allow terminals (perhaps via dial-up modems) to reach other hosts in the Internet.

Every Internet host is supposed to support a command called *whois* that can be used to look up directory information. Failing that, it is possible to telnet to SRI-NIC.ARPA and type WHOIS. That host is run by the NIC.

A great deal of information can be obtained by anonymous FTP (login anonymous, password guest) from SRI-NIC.ARPA, particularly from the files in the (NETINFO) directory. The specifications for most Internet protocols are on-line in documents called RFCs (Request for Comments) in the (RFC) directory ([83]; see especially RFC980 [41]). Most of the major Internet protocols have their actual specifications in *Military Standards for DoD Internet Protocols* [53]; most relevant RFCs and MIL-STD documents are collected and published in the 1985 *DDN Protocol Handbook* [82]. Many RFCs related to mail and domains have also been posted to the USENET newsgroup comp.doc.

ARPANET. Implementation of the ARPANET began in 1969 [15, 52]. ARPA administrators noticed that their contractors were tending to request the same resources (databases, powerful CPUs, graphics facilities, etc.), and decided to develop a network among the contractors that would allow them to share such resources. This network demonstrated the viability of long-haul packet-switched computer networks. It worked so well it had developed into a research utility (run by the Defense Communications Agency, or DCA) by the end of 1983, when it was split into MILNET, a production military network, and ARPANET, which reverted to research. In addition to the original goals of networking research and resource sharing, researchers almost immediately began using the network for collaboration through electronic mail and other services.

Policy is set by DARPA and executed by the Defense Data Network Project Management Office (see the "MILNET/MINET" section). The network is funded by DARPA and other government agencies. There were several versions of the early NCP (for network control protocols) and several early versions of TCP, but since 1983 the network has used the fourth version of the TCP/IP protocol suite above the PSN network layer (known as BBN 1822 after the report that describes it). Most of the links between PSNs are 56,000-bps leased lines. Response is quick

(except at peak load periods) and reliability is high. The name of the network is ARPANET; the name of the Internet is the ARPA Internet (*not* ARPANET Internet). There is also an Internet domain ARPA for hosts on the ARPANET and related networks, but that domain exists only to ease the transition to the real domains (COM, EDU, etc.), and will soon vanish. There are about 150 ARPANET hosts, all in the continental United States. Access to the ARPANET is officially limited to organizations doing research funded by federal money [17]. Most potential users will find it more productive to contact CSNET.

MILNET/MINET. MILNET is a long-haul military network that was built using the results of the ARPANET research. It split from the ARPANET in October 1983, but is still connected to the ARPANET by gateways at the internet layer. These gateways were originally intended to at least be able to limit traffic between the networks to mail only. Currently they pass all traffic as if the networks had not been divided (except for a performance penalty). Nonetheless, their PSNs form two disjoint sets, and the two networks could easily be separated if the need were to arise. More recently, the European nodes on MILNET have been separated into a network called MINET, which is also connected by gateways. Although MILNET eventually adopts most successful products of networking research done on the ARPANET, it does not usually participate directly in such research, since it is intended to be a stable operational network and service disruptions are kept to a minimum.

ARPANET, MILNET, and MINET are the main constituents of the Defense Data Network (DDN), which is a subset of the ARPA Internet and consists of networks that are directly managed by the Defense Data Network Program Management Office (DDN PMO), an Office of the Defense Communications Agency (DCA). Funding is by the U.S. Department of Defense (DoD). MILNET uses the same protocols as ARPANET, except MILNET has not yet adopted domain nameservers and still uses static host tables for host name to address mapping. MILNET PSNs are connected by 56,000-bps leased lines like ARPANET PSNs, but MINET hosts are connected by 9,600-bps links and reliability is lower. The gateways between MILNET and ARPANET are currently overloaded and form a severe bottleneck. Better gateway machines are expected to be installed soon. There are about 400 hosts, most in the continental United States, with some in Hawaii and Europe (the latter on MINET). There is a classified segment of MILNET in addition to the readily accessible part. Access is determined by the DoD.

There are other classified networks, including

some that still run NCP. The main DoD network uses switched TELETYPE messages. TCP/IP military networks other than MILNET and MINET in DDN are DISNET (Defense Investigative Network), SCINET (Sensitive Compartmented Information Network), and WINCS (WWMCCS Intercomputer Network Communication Subsystem—WWMCCS stands for World Wide Military Command and Control System). There are local networks in the Internet at military installations such as the Ballistics Research Laboratory (BRL).

Canada has the military network DRENET, which is an ARPANET-like PSN/TCP/IP network and is connected to the ARPA Internet.

Other ARPA Internet Networks. There are several unusual networks in the Internet. SATNET uses synchronous satellites to provide paths between the east and west coasts of the United States that are faster than the usual ARPANET or MILNET land lines, and there are satellite links to Hawaii and Norway. There are packet radio networks with nodes on mobile vehicles.

Many companies, schools, and government agencies have local networks that are part of the Internet. These include ethernets, token rings, broadband networks, and ARPANET-style PSN networks. Internet networks run the TCP/IP protocols on top of X.25 on public data networks. There are point-to-point connections over terminal lines, perchannel links, dial-up links, and T1 microwave links. These point-to-point links are usually used to connect higher speed networks. The speeds of

local networks may thus vary from 1,200 bps to 100 Mbps perchannel speeds or higher.

Many campus-sized organizations actually have several local networks. Since there is no need for people outside of the local organization to know the details of such internal networking arrangements, and since there is also a limit on the number of networks that the Internet core gateways can handle, many organizations arrange that their networks appear logically as a single network to the rest of the Internet, with subnets that are known only locally [60].

To be part of the Internet, a network must run the TCP/IP protocol suite and be connected (perhaps indirectly through another network) to one of the backbone long-haul Internet networks. Most are connected to either ARPANET or MILNET (but seldom to both).

CSNET

The purpose of CSNET is to facilitate research and advanced development in computer science by providing a means for increased collaboration among

those working in the field [12, 20]. The developers of CSNET noticed that electronic mail was the most popular service on the ARPANET. They proposed a network to provide electronic mail only and used it to connect institutions that did not have ARPANET access to those that did. CSNET is currently a logical network consisting of several physical networks, but serving a single community. All parts of CSNET are administered by a Coordination and Information Center (CIC) at BBN in Cambridge, Massachusetts [3]. The network has been self-supporting since 1985, though initial funding was provided by a grant from the NSF starting in 1981. Annual dues are collected from member organizations with rates set according to several classifications (usually either academic or industrial).

The only service supported on all the parts of CSNET is mail, transferred in ARPA Internet M822 format. But CSNET has in fact become a network built up of several parts that vary in their additional services, lower level protocols, and reliability, and other qualities. Some of these do support remote login, file transfer, and other services. With time, CSNET users began to realize that electronic mail alone was not enough. Though the ARPA syntax (e.g., user@host) is still used, CSNET is moving toward the new ARPA Internet domain name syntax.

The network is mostly confined to the United States and Canada, but has links to international sites in Australia, France, Germany, Israel, Japan, Korea, Sweden, and the United Kingdom. There are about 180 hosts in all, many of which serve as gateways to internal company networks or national networks. Perhaps thousands of hosts on such international networks can be reached through CSNET. Membership is limited to organizations "engaged in computer-related research or advanced development in science or engineering." Use of CSNET for commercial gain is explicitly prohibited.

CSNET CIC runs a nameserver, which is a direct database of CSNET users and sites. It can be accessed by mail or by remote login (from those components of CSNET that support remote login). The CSNET Info-Server allows retrieval of numerous documents by mail. To get a description of the information available and how to get it, mail a letter to cic@sh.cs.net with text containing the following

REQUEST: INFO
TOPIC: HELP
REQUEST: END

This information is also available over the ARPA Internet via anonymous FTP from [sh.cs.net](ftp://sh.cs.net).

CSNET CIC
cic@sh.cs.net
BBN Laboratories, Inc.
10 Moulton Street
Cambridge, MA 02238

Phonenet. Phonenet is a dial-up star network of about 128 hosts. A central relay computer (CSNET-RELAY) at CSNET CIC polls hosts on the network at mutually agreed upon times. Most such links are over 1,200-bps telephone lines, but 2,400-bps service has recently been installed, and some connections are over public X.25 links. The software used to manage the connections is called MMDF2, though many sites use PMDF, a Pascal subset of MMDF, combined with *sendmail* or the VMS mailer.

X25NET. X25NET uses X.25-based public data networks to support TCP/IP links. Because the TCP/IP protocol suite is used, additional services such as file transfer and remote login are provided. The use of that protocol suite also allows many of these hosts to be integrated into the ARPA Internet directly, if approved by DCA. There are about 18 hosts on X25NET.

ARPANET. About 25 hosts on the ARPANET proper pay CSNET dues and are thus logically considered part of CSNET as well as part of ARPANET and the ARPA Internet. However, all hosts in the ARPA Internet can be reached from any part of CSNET.

Cypress. Cypress is an attempt to provide ARPANET-like service to academic research departments at a cost as inexpensive as Phonenet service [42, p. 231]. It will use TCP/IP over 9,600-bps leased telephone lines. The nodes are called IMplets, and the first ones are small VAXes [50, p. 9]. Unlike Phonenet, Cypress will support file transfer, remote login, and other Internet services in addition to mail. Cypress will more resemble a tree network like BITNET than a star network like Phonenet, though there will probably be some redundant links.

MFENET

This network originated in the mid 1970s to allow access to a Cray 1 at Lawrence Livermore National Laboratories [42]. It has since grown, using several underlying network and transport protocols to support access to more supercomputers. The basic purpose of the network is to connect physics departments doing research in nuclear fusion, specifically in Magnetic Fusion Energy (MFE). It is funded and administered by the U.S. Department of Energy (DOE) and mostly managed from Lawrence Livermore.

Mail, file transfer, remote command execution,

and remote login are all supported on at least parts of the network. There are also specialized remote procedure calls for interactive graphics terminals. Some of the links use DECNET[®], while others use special-purpose protocols developed at Livermore. The use of nonstandard protocols has led to interoperability problems with other networks. Therefore, the DOE is considering moving MFENET to the TCP/IP protocol suite and perhaps eventually to the ISO protocols.

The existing links range from 9,600-bps to 56,000-bps leased lines to 122,000-bps satellite links. Speed between any two hosts depends greatly on the intervening links. Reliability is high. Addressing in DECNET networks is discussed in the sections on DEC's Easynet and SPAN. There are about 120 hosts on the network, all in the continental United States except for one in Japan. Four supercomputers are reachable: a Cray 1, a Cray X-MP/2, a Cray 2, and a Cyber 205. There are gateways to the ARPA Internet and possibly to other networks. Access is restricted to DOE-funded researchers.

SPAN

Planning for SPAN (the Space Physics Analysis Network) [28] began in 1980, and operations commenced in 1981. SPAN was originally oriented toward researchers in Solar Terrestrial and Interplanetary Physics, but is now expanding to serve other disciplines. SPAN is a multimission, correlative data comparison network serving projects and facilities of the American National Aeronautics and Space Administration (NASA) in collaboration with the European Space Agency (ESA). These agencies have traditionally set up data collection networks to serve specific space missions, but SPAN is mission independent, general purpose, low cost, and easy to connect to. (However, it is sometimes used to support specific missions, such as the ICE mission to the Giacobini-Zinner comet [75] and the encounter with Comet Halley [27]). It is an operational network in that it is not intended to promote the development of network technology, but it is a research network in that it provides an infrastructure for space-related research. It was not created in order to access supercomputers, but supercomputers are becoming more available through it.

Guidance for the network is provided by the users through the Data Systems Users Workers Group (DSUWG) and project scientists [29]. Direct administration is done by project managers, network managers, and routing center managers [30]. NASA pays for all the links, while other participating organizations pay for their own host computers and network

interfaces. Much of the original hardware, such as the routing center computers, came from NASA.

The upper layer protocols are DECNET (see the section on DEC's Easynet). The lower layers are provided by NASA's Program Support Communication Network (PSCN) and the NASA Packet-Switch System (NPSS). PSCN is a circuit switched network, that is, a collection of leased lines and microwave links. NPSS consists of X.25 links, some of them on public X.25 networks. The backbone of the network is four routing centers at the Goddard Space Flight Center in Greenbelt, Maryland, the Johnson Space Center in Houston, Texas, the Jet Propulsion Laboratory in Pasadena, California, and the Marshall Space Flight Center in Huntsville, Alabama. These are connected by 56,000-bps links. Each router serves as the center of a star of 9,000-bps links to the other institutions on the network. Reliability is becoming

DECNET addresses consist of 16 bits, 6 specifying an area and 10 specifying a node within the area. Since there are only 64 possible areas, management of area numbers is very important. Within DEC's DECNET-based company network, all area numbers are in use; thus direct gateways between Easynet and other DECNETs are problematic.

There are many DECNETs other than SPAN on the side of Easynet. They cooperate in assigning area numbers, with SPAN management providing a forum, especially for those networks interested in joining SPAN (ESA provides a similar forum in Europe). A major task of SPAN's routing centers is assignment of nodes to areas.

There are currently more than one hundred hosts connected directly to SPAN, all of them DEC machines. Outside of NASA, there are many participating universities and laboratories, such as the Los Alamos National Laboratory. There are many LANs indirectly connected to SPAN. Because other existing DECNETs want to join SPAN, the total number of hosts is expected to reach five hundred within a year. There is a transatlantic X.25 link between Marshall Space Flight Center in Huntsville, Alabama and ESA's Operations Centre (ESOC) in Darmstadt, West Germany. A 9,600-bps link was installed in September from Goddard to Germany, and one to Japan is expected by the end of the year.

SPAN can be reached from TELENET, and there are gateways to BITNET and the ARPA Internet. Access is limited to researchers in appropriate areas. For more information, contact

SPAN Project Manager
Code 633
NASA Goddard Space Flight Center
National Space Science Data Center
Greenbelt, MD 20771

DECNET is a trademark of Digital Equipment Corporation.

MAILNET

MAILNET originated as a joint project of the Massachusetts Institute of Technology (MIT), EDUCOM, and 15 pioneer sites, with some initial funding from The Carnegie Foundation. Unfortunately, the network is expected to vanish by the end of 1986 due to a lack of funds. MAILNET is an inexpensive mail network connecting heterogeneous computer systems at academic institutions. It is run by EDUCOM and is a star network around a Multics machine at MIT. Institutions with MAILNET hosts are charged an installation fee (\$2100) and a monthly service fee (\$190), plus usage charges based on the number and length of messages sent each month. Eighty percent of all MAILNET messages cost less than 20 cents. These fees pay for a high degree of support. Most mail transfers are done by telephone dial up from the central mail relay machine, though TELENET or BITNET can also be used. CSNET's MMDF software is used to coordinate the calls, and ARPANET SMTP protocols are used for addressing and transferring messages in RFC822 format. The only hardware required is a modem. Speed depends on the underlying transfer mechanism, but hosts are polled at least once a day. Reliability is high.

The old-style ARPANET (e.g., user@host) syntax is used. Gateways exist to the ARPA Internet, BITNET, CSNET, and JANET. Monthly traffic averages just over 12,000 messages from 1,800 users. There are about 30 hosts in the United States, Canada, and Europe. For more information, contact

Internet: postmaster@educom

MASTER%EDUCOM@WISCVM.WISC.EDU

master%educom.mailnet@mit-multics.arpa

The origins of JANET lie in the interconnection of centrally funded large university computer centers and research establishments in the United Kingdom. The first service desired was remote batch. The British Post Office Experimental Packet-Switching Service (EPSS) encouraged the development of general networking protocols in this community during the 1970s. A separate physical network was evolved to replace the point-to-point lines at the same time. The early network evolved into the Science Engineering Research Council Network (SERCnet) in 1977 and was renamed Joint Academic Network (or JANET) on April 1, 1984. JANET was established to provide network links to universities and research institutions in the United Kingdom and net access to the outside world.

Technical and administrative support is supplied

by the Network Executive (NE), based at the SERC Rutherford Appleton Laboratory, and by the Joint Network Team (JNT). JANET is funded by the Computer Board for Universities and Research Councils [81]; both JNT and NE are part of the Computer Board secretariat. No direct charges are made for usage. The annual Computer Board budget for JANET and university LANs is about 3.5 million pounds. (Not all university LANs are funded by this money.)

Local networks connected to JANET tend to be either X.25 campus switches, Cambridge Rings (CR82 standard), or ethernet (IEEE 802.3). The latter two are becoming increasingly popular. The long-haul network layer is X.25 over leased lines; some of these lines are digital. Higher layers are based on the Coloured Book protocol specifications:

Blue Book: File Transfer

Pink Book: Ethernet Protocols

Yellow Book: Network Independent Transport Service

Green Book: TS29 Terminal Protocol

Red Book: Job Transfer and Manipulation Protocol

Grey Book: Mail Protocol

Orange Book: Cambridge Ring 82, Hardware and Protocol Specifications

Fawn Book: Simple Screen Management Protocol

Development of these protocols started in 1979. They are sometimes called the Rainbow Book Protocols. Most long-distance links are 64,000-bps digital or 48,000-bps analog, and subscriber lines are mostly 9,600 bps.

JANET has a domain name system similar to that of the ARPA Internet, but the order of the domain name parts is opposite, with the root on the left. The system is centrally administered and in full use.

There are gateways to the British Telecom Packet Switchstream Service (BT/PSS) and from there to the International Packet-Switched Service (IPSS). There are also gateways to ARPANET (via University College London), EUnet (via The University of Kent), and EARN (via the Rutherford Appleton Laboratory). All the gateways have access controls because of funding considerations. Routing in the wide-area network is via X.25 addresses; this address space is independent of the CCITT X.121 space, but conforms to its requirements. Routing between the wide-area network and attached networks is via network level relays using extended addressing supported by the Yellow Book protocol.

Tabulations of the number of hosts can vary, depending on factors such as whether PADs are counted, whether those on local-area networks are included, and whether only registered ones are

counted or an attempt is made to estimate the actual number of connected hosts. There are about 915 registered hosts including those on local-area nets, but there are probably really about 1500 connected hosts. There are only about 20 hosts directly on JANET. For access, contact

Network Executive
Rutherford Appleton Laboratory
Chilton
Didcot
Oxfordshire, OX11 0QX
England

For the Coloured Book protocols, contact JNT at the same address.

EAN (X.400) Networks

There are several networks in Europe and elsewhere using the EAN implementation (first developed for Canada's CDNnet) of X.400 and other ISO protocols. They have an address format with a usual user presentation form that resembles ARPA Internet domains, but with an internal format of X.400 attribute lists [16, 45]. EARN is not an EAN network, but it will also migrate to X.400 (and other ISO protocols) by the end of 1987. This should be interesting, since BITNET in the United States is considering moving to the TCP/IP protocol suite (including RFC822 mail).

CDNnet. CDNnet is available to workers in the Canadian research, advanced development, and educational communities. It is autonomous of the Canadian Department of Defense. The first intermachine messages were exchanged on CDNnet in 1983 [43, 62, 87]. The network currently uses the EAN implementation of X.400, although other X.400 software may be approved as it becomes available. EAN was and is developed at the University of British Columbia. Work was begun late in 1981 and tracked the standards work already in progress by IFIP and CCITT. Sydney Development has the commercial rights to EAN. The primary purpose of CDNnet is resource sharing and collaboration among researchers.

The network is administered by CDNnet Headquarters at the University of British Columbia, and by representatives at each member organization. CDNnet policies are set by a Steering Committee representing its users. EAN development work has been supported by grants from the Canadian Natural Sciences and Engineering Research Council (NSERC) since November 1981. The present NSERC grant is under its university-industry program; it is a three-

year grant that will terminate in 1988 [25]. The grant includes support for ongoing research and development, for the initial support of CDNnet, and for cooperation with a commercial organization (Sydney Development Corporation). CDNnet is to be self-sufficient by 1988. Annual dues will be collected, starting in 1987 with rates set according to the type of organization (e.g., educational, government, non-profit, commercial) and according to the number of CDNnet hosts at an organization. Organizations pay the telecommunications costs for connections to other organizations. CDNnet does not have usage charges except to recover the costs associated with gateways and bridges to other networks.

The EAN implementation conforms to CCITT ISO specifications at the session (CCITT X.225, ISO 8327), transport (TP0: CCITT X.224, ISO 8072, and 8073), and network (X.25, PSTN, DECNET, etc) layers. TTXP is also used as a network layer, based on the specifications of CSNET's MMDF.

Mail is the basic application service provided. X.400 also provides receipt notification, which is widely used in CDNnet and in the other EAN networks. This is implemented in EAN as follows: the sender requests this service, a receipt report is returned to the sender when the recipient displays the body of the message. USENET news is available on at least part of the network. There is a directory service for locating users of CDNnet and other EAN-based networks. Remote login is available to hosts with X.25 service. EAN implementations exist for 4.2BSD, VMS, System V, and will soon be developed for VM/CMS.

Most long-haul links are X.25 at 2,400 bps, though they vary from 1,200 bps to 9,600 bps (the range offered by the Canadian PDN Datapac). There are some leased lines. Mail delivery is usually accomplished within minutes, and reliability is high.

The CDNET address format is *user@subdomain.cd*, where *subdomain* is composed of a list of one or more simple names separated by dots. This is actually the format for presentation to the users. Internally, addresses are represented in binary form as X.400 Originator/Recipient names; the exact mapping is likely to change with the next version of X.400, in particular because of work on directory services, though it is possible that the same presentation format will continue to be used. For messages from the outside world, the addressing is more complicated (see Table III).

CDNnet has about 65 hosts, the busiest (*ean.ubc.cdn*) processing about 2000 messages a day. There are gateways to CSNET, MAILNET, BITNET, and UUCP, and close connections to the other EAN

networks. ARPANET can be reached indirectly through CSNET, and EUnet through either UUCP or the European EAN networks. Access is restricted to organizations engaged in research, advanced development, and education. The network may not be used for commercial purposes.

For access, contact

CDNnet Headquarters
cdnnet-hq@ean.ubc.cdn
Computing Centre
University of British Columbia
Vancouver, British Columbia
Canada V6T 1W5

European EAN Networks. The EAN protocols have spread so rapidly into new nations that it is interesting to track their progress (see Table IV). There are actually no EAN networks as such in the United Kingdom or Australia: There are merely gateways to the national networks, much as Australia has UUCP gateways. The German network is not strictly EAN-based. CDNnet, the progenitor of all the others, will not remain solely EAN as other X.400 implementations become available.

The objective of the EAN networks is to establish communication links for the European research community, in cooperation with RARE (p. 000). The networks use the EAN implementation of X.400. It is expected that they will migrate to other implementations in a few years due to the preliminary nature and lack of support for the current implementation. Most of the sites connected inside the same network are linked by 9,600-bps leased lines. Inter-network links consist mostly of X.25 public switched networks, with some using 9,600-bps lines. Mail de-

livery in minutes to hours is usual, with medium reliability.

Naming, addressing, and routing are the same as for CDNnet. Methods for reaching EAN networks from non-EAN networks vary greatly depending on the network of origin and the locations of the sender and the addressee. There is a gateway at CERN that connects EARN, UUCP, and the EAN networks. For access, contact

Alf Hansen
skjesol@vax.runit.unit.uninett
alf_hansen%vax.runit.unit.uninett@nta-vax.arpa
Trond Skjesol/Alf Hansen
N-7043 Trondheim-NTH
Norway

SMARTIX/COSAC

COSAC (Communications SAns Connections) is a French research network. Development work started in 1981 at the Centre National d'Études des Télécommunications (CNET). Version 3 has been operational since 1984 and has some restrictions. Version 5 is in development now and will be a full X.400 implementation, to be operational by the end of 1986. COSAC is administered by CNET, which also funds the network through CNET, INRIA (Institut National pour le Recherche d'Informatique et l'Automatique), CNRS (Centre National de la Recherche Scientifique), and Bull.

SMARTIX is intended as a generalization of the ideas of COSAC. It will use an implementation of X.400 by INRIA. Funding for SMARTIX will come from the French government and will involve CNET, INRIA, ADI (Agence de l'Informatique), Bull, and CNRS.

TABLE IV. Timetable for the Development of the European EAN Networks

Country	Version	Gateway	Date
Canada	CDNnet	EDN	Apr. 1983
Norway	UNINETT	UNINETT	Mar. 1, 1984
Switzerland	CERN	CERN	Oct. 7, 1984
UK	UK	UK	Nov. 2, 1984
Sweden	SUNET	SUNET	Nov. 27, 1984
Switzerland	CHUNET	CHUNET	Dec. 9, 1984
Germany	DFN	DFN	June 17, 1985
Ireland	IRL	IRL	Aug. 22, 1985
Italy	OSIRIDE	I	Nov. 12, 1985
Spain	IRIS	E	Dec. 3, 1985
Australia	AU	AD	Dec. 23, 1985
Netherlands	NL	NL	Mar. 17, 1986

† First intermachine message.
‡ First network established.
Table courtesy of John Demco.

COSAC uses the CCITT X.400 protocols over X.25, with ISO transport and session protocols plus FTAM for file transfer. Local links use 64,000-bps X.25 links, and long-distance ones use TRANSPAC, the French public data network. There is a gateway with FNET (the French UUCP network; part of EUnet). It is possible to get to the ARPA Internet, CSNET, and BITNET from COSAC. Between CSNET and COSAC, the French CSNET host france.csnet is used. COSAC has 27 hosts in France, of which about a dozen each are Multics and UNIX machines, a couple each are IBMs and DEC-20s, and one is a VMS VAX. Most are in the environs of Paris or provincial capitals, though the two DEC-20s are actually in Dublin, Ireland. For access, write to

mulcnet/kintzig@france.csnet (Claude Kintzig)
seismo!mcvax!inria!gipsy!ch (Christian Huitema)
adi/bernard@france.csnet (Jean-Luc Bernard)

or to Claude Kintzig at

Projet SMARTIX
CNET (PAA/TIM)
38-40 Rue General Leclerc
F-92131 Issy-les-Moulineaux
France

DFN

DFN (Deutsche Forschungnetz) is the national research network in West Germany. There was early German networking activity at the Hahn Meitner Institut in Berlin, where an X.25-based network called HMI-net was developed. There was also an academic network between two universities there. The largest research network in the country is BERNET, which is still in Berlin. In 1982 there was a move to expand BERNET to be a Northern German network. However, a study conducted by Stanford University recommended as an alternative a national network to provide ARPANET-like services. DFN was started to implement this idea.

The purpose of DFN is to develop protocols and implementations in the ISO suite that can be used for resource sharing and collaboration among researchers nationwide and communications with foreign researchers. DFN is the German part of RARE. The West German Ministry of Research and Technology funds DFN and has about 15-20 people working on the DFN project, though all implementation work is contracted out.

DFN uses X.400 for mail, plus file transfer and remote job entry using protocols designed for the network but compatible with the OSI suite. The EAN implementation of X.400 is used currently, though a German implementation is being devel-

oped. The network layer is X.25, which supports remote login. Most links are 9,600 bps.

There are about half a dozen existing hosts supporting mail, all in the Federal Republic of Germany. Immediate plans are to have about 30 hosts: 10 4.2BSD, 10 System V, and 10 VMS. Gateways exist to EUnet, EARN, CSNET, and the EAN network. For access, contact

DFN-Verein
Pariser Strasse 44
1000 Berlin 15
Bundesrepublik Deutschland

There is also a newsletter published by a consortium of people from several German networks, giving a comprehensive view of all networking activities in Germany:

Neueste Netz Nachrichten
Universitat Karlsruhe, Zirkel 2
D-7500 Karlsruhe 1
Bundesrepublik Deutschland

ROSE

ROSE (Research Open Systems for Europe) is the principal development project of IES, the European Information Exchange System [4]. ESPRIT is the European Strategic Program in Information Technology, the European Economic Commission (EEC) started on the ROSE implementations in 1982. The goal of providing an infrastructure for cooperative research and development projects with ESPRIT, and eventually for other projects of

kinds in Europe. It is also a proving ground for the use of the ISO protocols in an environment of heterogeneous machines and both wide- and local-area networks. Funding comes from the EEC and goes to five industrial partners who do the work: Bull of France, GEC and ICL of the United Kingdom, Olivetti of Italy, and Siemens of West Germany. Some tasks are subcontracted.

Services eventually expected under ROSE include mail, conferencing, file transfer (including text files), remote command execution, and remote login. The UNIX operating system has been chosen as the first implementation system. Initially, existing implementations of protocols already in widespread use on UNIX, such as UUCP, will be used (see also the "EUnet" section). Eventually, all protocols will be those of the ISO suite. Those protocols are chosen from those recommended by the Standards Promotion and Application Group (SPAG) [80], which is a group of a dozen European manufacturers interested in promoting common networking standards in their products.

Remote terminal access will be accomplished by X.3, X.23, and X.29 PADs, file transfer will be by ISO 8571 (FTAM), mail will use the X.400 series of CCITT recommendations, session will be implemented as ISO 8326 and ISO 8327, transport will be ISO 8072 with classes 0, 2, and 3 over X.25 and class 4 over CSMA/CD networks, the internet layer will be ISO 8473, and the network layer will mostly be X.25 and X.75.

The end-to-end addressing convention to be used in ROSE with the ISO protocols is a three-level hierarchy of eight octets for the name of the remote network, eight octets for the system on the LAN, and two octets for the transport selector. This allows gateways between networks to be the only machines that need to know about the interconnection topology of the networks. The transport selector could allow the user to choose UUCP instead of the ISO session service. A prototype network is just being set

COMPANY NETWORKS

Xerox Internet

Xerox was a pioneer in network research and in fact (ad. Ethernet [91] along with the Xerox Palo Alto Research Center (PARC) Universal Packet protocol [5] and the Xerox network system [92], protocol suites. An internetwork among company sites, the Xerox Research Internet was developed by 1976. The Corporate Internet split from the RIN about mid 1985.

CIN and RIN are highly interconnected and form the Xerox Internet (XIN). The CIN was a stable backbone network for various needs, and the RIN is primarily intended for research and development.

Administration of the networks is distributed among several groups within Xerox. RIN is administered mostly by Xerox PARC, while CIN is mostly run by other divisions within the Xerox Corporation. Both RIN and CIN are funded by the Xerox Corporation.

Several protocols are used by XIN. The XNS protocol is used by both CIN and RIN. RIN also uses PUP as well as TCP/IP on some of its nets. Reliability is reported to be at least as good as that of ARPANET. Transfer speed on RIN is maintained in part by 100-bps leased lines and T1 Microwave links. Naming and addressing are handled differently for internal and external users. The mail system in RIN is called Grapevine, and the name system in CIN is called Clearinghouse. For internal names, the following examples are typical:

Grapevine name: User.registry (e.g., JLarson.PA, PA = Palo Alto)

Clearinghouse name: Name:Domain:Org (e.g., John Doe:OSBU North:Xerox)

XIN communicates with the outside world via two hosts on the ARPA Internet: Xerox.COM and parcvax.Xerox.COM. Xerox.COM (formerly Xerox.ARPA) is the ARPANET-Grapevine mail gateway, and parcvax.Xerox.COM is used for telnet and FTP. Several mail gateways connect Grapevine to the CIN mail system. Example addresses (from the ARPA Internet) are

(to RIN/Grapevine): User.registry@Xerox.COM
(e.g., JLarson.PA@Xerox.COM)

and

(to CIN/Clearinghouse):
Name.foreignRegistry@Xerox.COM
(e.g., JDoe.osbunorth@Xerox.COM)

Certain aliases are maintained at the Xerox.COM gateway for ease of external addressing (e.g., Postmaster@Xerox.COM goes to Postmaster.pa).

The services offered by RIN and CIN include remote login, file transfer, mail, remote procedure call, distributed file system, distributed computation, and many others. Both CIN and RIN support many thousands of machines. Grapevine, a very large distributed mail and name system for RIN, supports about 4000 users around the world. CIN has its own distributed mail system and about 8000 users. The XIN is international in scope, maintaining links to sites in Japan, England, and Canada, as well as to numerous sites within the United States. For access, contact

Postmaster@Xerox.COM.

DEC's Easynet

DEC maintains an internal engineering network called Easynet. DEC began its network endeavors as one of the pioneers of the ARPANET. The company went on to develop its own network software, called DECNET, and to make it available to DEC customers by 1976. Easynet was started in 1978.

The basic network capabilities provided by DECNET include intersystem file access and transfer, electronic mail, intersystem resource sharing, interprocess communications, adaptive routing, and remote login. These are enhanced by Easynet services such as ELF (a DEC Employee Locator Facility), Videotext Infobases, on-line network conference discussions (interactive bulletin boards), system monitoring, and communications for general operational issues.

Easynet is administered by DEC Corporate Telecommunications with funding provided by DEC. Easynet uses the DECNET protocol. Reliability is reported to be at least as good as the ARPANET's. Easynet's speed is maintained by 10-Mbit ethernet and 56,000-64,000-bps backbone intersite links, with lower speed links to sites with lower traffic requirements. Addressing is the same as for the Internet: for instance,

user%host.DEC@decwrl.DEC.COM

For UUCP, the correct address mode is

{ucbvax, decvax, ...}!decwrl!enetnode.dec.com!user

Easynet has more than 10,000 hosts. Assuming that two-thirds of DEC's employees are users, there are about 60,000 network users. The network is international in scope, extending throughout North America, the Caribbean, Europe, the Near East, Australia, New Zealand, and the Far East.

The gateway on the DEC end is decwrl.dec.com, a VAX 11/750 running 4.2BSD UNIX, along with RHEA, an Easynet node in the same room that is another VAX 11/750 running VAX/VMS plus Eunice with TCP/IP. These two machines share the responsibility for gatewaying mail between Easynet and the ARPA Internet, CSNET, and UUCP.

For access, contact

johnsson%covert.DEC@DECWRL.DEC.COM

Digital Equipment Corporation
Distributed Systems Architecture
550 King Street
LKG1-2/A19
Littleton, MA 01460-1289

IBM's VNET

IBM has several internal networks supporting mail, remote login, and file transfer for company operations. The main internal network is actually two distinct networks that together form what is generally called VNET. The RSCS (for Remote Spooling and Communications Subsystem) network [31] comprises the mail and file transfer part of VNET and currently has approximately 2200 nodes. The PVM (Passthru VM) network provides a remote login facility for VNET. There are about 1100 PVM nodes. VNET RSCS and PVM nodes are found in North and South America, Africa, Middle East, Europe, Australia, and Asia. As mentioned, there are other internal networks; for example, VIBTS (VTAM Integrated Bulk Data Transfer System) is a fast network consisting mostly of T1 microwave links. VIBTS is for transferring memory images during debugging and other activities that require fast access to huge

files. Another internal network is CCDN, which is used solely for remote login.

VNET and the RSCS protocols began in 1972 as an ad hoc project of some IBM employees who felt that the available alternatives did not meet their needs. IBM eventually adopted the new network, and in this sense, VNET is the UUCP of IBM. VNET is IBM's internal network, providing services such as mail, remote login, and file transfer for company employees. VNET administration is run by a group called the VNET Project Team, which was formed in 1978. This team maintains and sets network guidelines. There is also a VNET corporate office that was established in 1982. VNET is funded by the IBM Corporation.

The RSCS protocols are used for mail transfer. These are the same as the ones BITNET uses. VNET links are typically 9,600-bps leased lines, though they vary from 2,400-bps to T1 speeds. Addressing is similar to the Internet style, for instance,

From ARPA Internet: USERID@IBM.COM
VNET to VNET: USERID at HOSTNAME

This will result in either a direct transmission of mail to the IBM employee or in the generation of a message that the intended recipient needs to enter.

Some non-IBM mailing lists (including the USENET newsgroups) are gatewayed into IBM conferencing systems. VNET has grown from a few hundred hosts a few years ago to about 2200 now. The VNET PVM nodes are perhaps not a subset of the VNET RSCS nodes, but for all practical purposes, they can be considered as such. VNET gateways to both CSNET and BITNET that operate within the restrictions stated below. Through the gateways it is possible to send mail to any registered IBM employee. File transfer and remote login are supported for outsiders. The CSNET gateway can support mail to any registered IBM employee on an internal machine. The BITNET gateway at Yorktown Heights is not as flexible. It can support mail only to those hosts at Yorktown Heights and Almaden (not that the IBM Research Division site on the West Coast has moved from San Jose to Almaden and is now known as IBM Almaden Research, often referred to as ARC) that are in the CUNY BITNET (RSCS) node table. There are also other BITNET gateways that are more flexible, and there is a report that describes them [88].

IBM maintains its own internal security to provide limited access to VNET. For a person within IBM to be able to send or receive mail, it is necessary to obtain an account on a VNET node; in addition, to

exchange mail with external networks it is necessary to register. Most professionals within IBM are VNET users, although a relatively small number are registered to talk to the outside world via gateways. An outsider need not be registered to send mail to VNET, but the person receiving the mail inside IBM must be registered. Many of the people in the Research Division (Yorktown Heights, Almaden, and Zurich) and the various Scientific Centers are registered, although most others in IBM are not. For additional information, contact

postmaster@IBM.COM

AT&T

AT&T has some internal networks, most of which are internally developed transport mechanisms. Their most widely used networks are UUCP and USENET, which are not limited only to that corporation and which are discussed later. All internal AT&T networks support UUCP-style h1!h2!h!u source routing syntax and thus appear to the user to be UUCP. Within AT&T, UUCP links are typically over 1,200-bps dial-up telephone lines or Datakit (see below).

Among AT&T's other networks, CORNET is an internal analog phone network used by UUCP and modems as an alternative to Direct Distance Dialing (DDD). Datakit is a circuit-switched digital net and is similar to X.25 in some ways. Most of Bell Laboratories are trunked together on Datakit. On top of the Datakit service, people run UUCP for mail and remote login. In addition to host-to-host connections, Datakit supports RS232 connections for terminals, printers, and hosts. ISN is the version of UUCP supported by AT&T Information Systems. The Bell Laboratories in Holmdel, New Jersey, uses ISN for internal data communication. BLICN (Bell Labs Internal Computing Network) is an IBM mainframe network dating from the early 1970s. Programmer's Workbench (PWB) was a common version of the UNIX operating system. Many Bell Laboratories machines with PWB-style RJE links use BLICN for mail and netnews for other UNIX machines. A major USENET host uses this mechanism to connect to about 60 neighbor hosts. BLICN connects Bell Laboratories installations in New Jersey, Columbus, Ohio, and Chicago, and links most computer center machines. BLN (Bell Labs Network) is a T1 Hyperchannel at Indian Hill, Chicago. BITNET Internet is a TCP/IP internet. It is not a Bell AT&T network, though some of the best machines are on it. There are many ether-nets connected by TCP/IP over Datakit. This internet may soon be connected to the ARPA Internet.

ACCUNET is AT&T's commercial X.25 network. AT&T MAIL is a commercial service that is heavily used within AT&T Information Systems for corporate internal mail.

COOPERATIVE NETWORKS

BITNET

BITNET (Because It's Time NETwork) [24] is a cooperative network serving over 1300 hosts located at several hundred sites (mostly universities) in 21 countries, as shown in Table V, on the next page. BITNET is a communications link between universities and research centers with few requirements or restrictions other than that a site must acquire a leased line to facilitate the connection to another BITNET node and, in the spirit of a cooperative network, be willing to serve as a connection node for at least one new member. (Commercial members cannot currently communicate with other commercial members, however.) This concept of virtually unrestricted access and the absence of membership fees not only characterizes the cooperative essence of BITNET, but distinguishes it from other networks available for interinstitutional communication such as CSNET. Any university or college that is able to connect to BITNET possesses unlimited collaborative possibilities for both academic and administrative purposes for faculty, staff, and students. This policy is similar to that of UUCP and USENET, except that it is more limited to academic institutions. BITNET also resembles CSNET somewhat, but is not as centralized and has never been supported by the government.

There are three main constituents of the network: BITNET in the United States and Mexico, NET-NORTH in Canada, and EARN in Europe. There is also AsiaNet in Japan, and there are plans for expansion into South America. The distinctions are purely political, and mail can be freely exchanged between any two hosts.

BITNET in the United States. BITNET had its beginnings in 1981 when the first two sites, City University of New York (CUNY) and Yale University, were connected on May 5. BITNET was originally used primarily for collaboration and communications among systems programmers at university computation centers. Since existing IBM networking software was used, BITNET was initially a network of IBM hosts. As BITNET has grown, software to emulate the protocols has been developed by commercial vendors and members of the BITNET community.

The basic premise for establishing BITNET was to provide a communications network among universi-

TABLE V. BITNET Hosts as of May 1, 1986

	Country	Abbreviations	Hosts
BITNET	U.S.A.	U.S.A.	844
	Mexico	MEX	1
Totals	2		845
NETNORTH	Canada	CAN	91
Asianet	Japan	JPN	7
EARN	Austria	A	6
	Belgium	B	13
	Israel	IL	38
	Switzerland	CH	22
	Germany	D	130
	Denmark	DK	13
	Italy	I	31
	Spain	E	8
	France	F	39
	Netherlands	NL	39
	Finland	SF	7
	Greece	GR	2
	Ireland	IRL	4
	Norway	N	1
	Portugal	P	1
	Sweden	S	6
	U.K.	GB	9
Totals	17		366
All parts of BITNET	21		1211

Table courtesy of Henry Nussbacher.

ties with no special requirements, restrictions, or fees for membership. Today BITNET is used by scholars and administrators from a variety of different disciplines. Services provided include electronic mail (RFC822), file transfer, and interactive messages. The interactive messages allow several users to communicate interactively while experiencing only moderate delays, usually less than eight seconds.

User, technical, and administrative support is provided by the BITNET Network Support Center, which is operated jointly by EDUCOM and CUNY. EDUCOM handles the user services and administrative support through the BITNET Network Information Center (BITNIC), the purpose of which is to promote the use of BITNET in higher education. Services provided include an on-line directory, paper and electronic newsletters, end-user documentation, workshops, seminars, and conference presentations. In addition to these direct user-service functions, BITNIC also provides administrative support by negotiating for software and equipment discounts and by archiving network procedures and policies. CUNY directs the technical support, systems maintenance, and software development efforts via its BITNET Development and Operations Center (BITDOC). BITDOC maintains the Support Center computer and is dedicated to improving existing

BITNET services and implementing new ones. Through this coordinated effort, BITNIC and BITDOC are able to provide a very high level of support for BITNET. BITSERVE is a help facility organized and operated by BITDOC at CUNY that offers a BITNET news service, a user directory, and a list of BITNET sites and computers, and is being expanded to include information on conferences, software, and special facilities available to BITNET members. BITSERVE is presently accessible to all BITNET sites and is being enhanced to allow access from other networks.

Considering the absence of membership fees, the cost to institutions for BITNET is small and the restrictions are minimal, compared to other networks. The site must provide the connection to an adjacent site and either acquire IBM's VM-based RSCS from a vendor or obtain emulation software free of charge from BITDOC. Government support provides the funding for the support sites. IBM has provided funds for BITNIC, but that funding ends soon. However, EDUCOM and the BITNET board of directors realize that BITNIC is essential to the success and efficiency of the network. As of now no decision has been made on how to support BITNIC, but it is clear that a member fee will have to be initiated for BITNIC to remain in existence.

BITNET started as a network of IBM hosts, and most BITNET hosts communicate using an IBM communication environment. BITNET uses the NJE (Network Job Entry) protocol. Most hosts use RSCS. Today there are many non-IBM hosts using emulation software to provide the appropriate protocols for the DEC, VMS, UNIX, and Sperry environments. Hosts are interconnected by leased phone lines supporting 9,600-bps data transmission. The mail and file delivery delay ranges from minutes to hours with medium reliability. An unusual feature of BITNET is that there is usually exactly one path between any two hosts. A geographic map of the network shows it to be a tree network rooted at CUNY. This means that nodes near the root of the tree are more important. An example is CUNYVM. If this node goes down, the network will still run, but will essentially be split into two networks, with files and mail queued at the nodes nearest CUNYVM. The addressing format from the ARPANET to BITNET is

host.BITNET@WISCVM.WISC.EDU

Methods for specifying addresses to the mail system vary widely, depending on the host. There are currently over 800 nodes distributed among several hundred sites. The scope is the United States plus Mexico, although there are direct links to NORTH and EARN. Gateways exist between BITNET and MAILNET, EDUCOM, CSNET, and ARPANET, and there is also restricted access to BITNET. For access, contact

network information center (BITNIC).
INFO@BITNIC

Internet:

%BITNIC.BITNET@WISCVM.WISC.EDU

NETNORTH. NETNORTH provides communication for a number of Canadian academic and research sites, and was designed using the same techniques and several of the same basic assumptions as BITNET. The network currently has over 90 nodes, and direct links exist to BITNET and EARN. NETNORTH and BITNET are connected by a leased line between Cornell University and the University of Guelph, in Ontario. Plans are also underway to provide connections between NETNORTH and other Canadian networks.

EARN. The European Academic Research Network, links over 150 hosts at over 100 institutions in 18 countries. EARN is similar to NETNORTH in that it is based on the same design principles and philosophy as BITNET. EARN is an integral part of BITNET, which means that all Euro-

pean nodes must be listed at all U.S. BITNET sites and vice versa. The central administrative and technical services are handled by one central computer in each country (analogous to CUNYVM for BITNET). EARN has an administrative branch very similar to BITNIC. IBM supports EARN and its gateway to BITNET. For example, IBM will fund BITNET links to several major European EARN sites until April 1987. Connections to other sites are welcome, but the connecting host must pay for such new links.

The gateway between EARN and other nets is the same as for BITNET. For example, to send mail from the ARPA Internet to EARN, the address format would be

user%EARNhost.BITNET@WISCVM.WISC.EDU.

There are more than 350 hosts in 17 European countries. Gateways to several national academic networks in Europe are planned. EARN is currently expanding into Turkey, Iceland, and Portugal.

Country coordinators are available for each of the 18 countries currently connected to EARN. The large majority of nodes have installed a mailbox named INFO. Users on other networks may use the internet address,

INFO%EARN_NODE.BITNET@WISCVM.WISC.EDU.

FidoNet

FidoNet uses a telecommunications package for personal computers that was developed by Tom Jennings in 1983. FidoNet is basically an extension of the Fido Bulletin Board System (BBS), which provides the electronic mail portion and makes the Fido BBS unique in that unattended mail transfer between other "nets" and their "nodes" is not often offered with bulletin-board services.

FidoNet provides a network for personal-computer users in the spirit of BITNET's utility to the academic community. There are no special requirements or fees for membership. The software is shareware. There is no distribution cost, though users are asked to donate a small sum. The network is distributed without sources. The main services are electronic mail for personal computers and access to USENET newsgroups for personal-computer users. Since FidoNet is part of the Fido BBS, all of the services of the BBS are also available (these are covered in the section on bulletin-board systems).

The administrative node of FidoNet is located in St. Louis. The USENET/Fido gateway is administered by Bob Hartman at FidoNode 101/101, vaxine!sparks!m!n!user. FidoNet is funded by its users, who must send the Sysop (System Operator) of the Fido node a cash retainer (usually \$5 or \$10) to

cover the phone costs when a mail message is sent. Since FidoNet operates late in the evening (after 11 PM local time) when phone rates are lowest, the cost per message is minimal. Sysops usually provide the machines (since Fido and FidoNet only run under MS or PC DOS—the operating systems for IBM PCs and compatibles—the machines are IBM PCs and compatible machines) and donate their time at no cost. FidoNet is not designed to be a commercial venture.

FidoNet uses the Fido protocol with connections made at 1,200 or 2,400 bps. A FidoNet address is composed of the user's name, a net (a region or host), and a node (a Fido) from the available list; for instance,

user Net_number Node node_number.

The bulk of the nodes are in the United States and tend to be clumped in metropolitan areas such as St. Louis, Boston, and Chicago. However, there are a number of nodes in Europe, a few in Indonesia, and even one in Alaska. Fido BBS is installed on over 500 hosts (IBM PCs or compatibles), and new hosts or nodes are joining on a regular basis. To access FidoNet, contact the local IBM PC User Group to see if there is a local Fido node in your area. If so, they should have the dial-up number. For more complete information about FidoNet in general, contact

Tom Jennings

FidoNet: Fido's_Board Net 125 Node 1

UUCP:

ihnp4!encore!vaxine!spark!125!1!tom_jennings

ACSNET

ACSNET (Australian Computer Science Network) is the main network in Australia and is based on the Sydney UNIX Network (SUN) software developed at the University of Sydney [22]. The network started in 1979 and connected a machine at Sydney to another at the University of New South Wales. It currently spans the continent and is closely connected to networks elsewhere [47]. The purpose of the network is to support mail traffic and file transfer among research, academic, and industry users. The underlying transport protocols are also used to support the USENET news network in Australia. There is no central administration, though this may change in the future. At present the original developers and the international gateway operator act as coordinators. There is no government funding: Each host pays for its own links.

The original protocols were called SUN-I and supported remote login, file transfer, and multiplexed

protocols. Dynamic routing was added in 1980, but only applied to mail and file transfer. SUN-II was similar, but allowed intermittent (dial-up) links as well as dedicated ones, plus a method of layering SUN-II on top of other networks (such as CSIRONET).

The current version is a complete redesign and reimplemention done in 1983 and is called SUN-III [48]. It is layered in the traditional networking manner and provides a message delivery service with implicit (system) routing and domains in order to support higher level protocols including file transfer, electronic mail, news, remote printing, simple directory service, and a number of experimental services. It can transfer messages in both directions simultaneously over full-duplex links. It supports multiplexing, which is useful with USENET news and also with mail addressed to users on multiple hosts.

The transport protocol can make use of any form of virtual circuit between hosts. The links currently in use include leased lines, dial-up lines, X.25, and CSIRONET. CSIRONET is a government research network originally developed to connect terminal users in remote areas to a central facility; it now provides virtual circuits between host machines. CSIRONET is slow but cheap for long distances, which there are many in Australia. There is a plan to migrate the system to X.400 in the next few years. Most links run at 1,200 bps, and reliability is high.

ACSNET has a domain naming syntax [49] similar to that for ARPA Internet domains. The domain OZ.AU is registered with the Internet and cannot be interpreted as a subdomain, OZ, for ACSNET, or a country domain, AU, for Australia. There are currently no other subdomains of AU, but there are subdomains within OZ.AU, many of which are for distributed organizations. Domains are used for routing in ACSNET, so connections between machines determine domains more than anything else. Hosts can register in any subdomain. In practice, this means major hosts are directly in OZ, and everything else is in subdomains.

There are several UUCP gateways to North America and Europe, all from Melbourne and all using X.25. There is a CSNET link to the United States. There are EAN links to Canada, the United Kingdom, Germany, Norway, and Switzerland. In addition, there are several 1,200-bps dial-up links to North America (New Hampshire, New Jersey, and California). There are about 300 hosts throughout the Australian continent on ACSNET.

For general information, contact

postmaster@munnari.oz.au
seismo@munnari!postmaster

For connection requests, contact

acsnet-request@basser.oz.au

seismo!muninari!basser.oz!acsnet-request

or the ACSNET Coordinator at

Department of Computer Science
University of Sydney
New South Wales, 2006
Australia

UUCP

The name "UUCP," for UNIX to UNIX CoPy, originally applied to a transport service used over dial-ups between adjacent systems [63]. File transfer and remote command execution were the original intent and main use of UUCP. There was an assumption that any pair of communicating machines had direct dial-up links, that is, that no relaying was done through intermediate machines. By the end of 1978, there were 82 hosts within Bell Laboratories connected by UUCP. Though remote command execution and file transfer were heavily used, there is no mention of mail in the standard reference [64]. There was another similar network of "operational" links with UUCP links that were apparently outside Bell Laboratories, but still within the Bell System. The two networks intersected at one Bell Laboratory machine.

These early networks differed from the UUCP network in assuming direct connections between communicating hosts and in not having a mail service. The UUCP mail network proper grew from the early networks and spread as the programs were distributed as part of the system.

Remote command execution can be made to work over successive links by arranging for each job in turn to submit the next one. There are several programs that do this: Unfortunately, they are all incompatible. There is no facility at the transport layer for routing beyond adjacent systems or for error acknowledgment. All routing and end-to-end reliability support is done explicitly by application programs implemented using the remote command execution facility. There has never been any remote mail facility associated with UUCP, though the *cu* dial-up programs are sometimes used over the same telephone links.

The UUCP mail network connects a very diverse set of machines and users. Most of the host machines run the UNIX operating system [73, 74]. Mail is the only service provided throughout the network. In addition to the usual uses of mail, much traffic is

generated as responses to USENET news. The same underlying UUCP transport mechanisms are also used to support much of USENET.

The UUCP mail network has many problems with routing (it is one of the few major networks that uses source routing) and with its scale. Nonetheless, it is extremely popular and still growing rapidly. This is attributable to three circumstances: ease of connection, low cost, and its close relationship with the USENET news network.

Mailing lists similar to those long current on the ARPANET have recently increased in popularity on the UUCP mail network. These permit a feature that USENET newsgroups cannot readily supply: a limitation of access on a per-person basis rather than on a per-host basis. Also, for low-traffic discussions mailing lists are more economical, since traffic can be directed to individuals according to their specific interests.

There is no central administration. To connect to the network, one need only find one machine that will agree to be a neighbor. For people at other hosts to be able to find your host, however, it is good to be registered in the UUCP map, which is kept by the group of volunteers known as the UUCP Project [84]. The map is posted monthly in the USENET newsgroup *comp.mail.maps*. There is a directory of personal addresses on the UUCP network [44], although this is a commercial venture unrelated to the UUCP Project.

Each host pays for its own links; some hosts encourage others to connect to them in order to shorten mail delivery paths.

There is no clear distinction between transport and network layers in UUCP, and there is nothing resembling an Internet Protocol. The details of the transport protocol are undocumented (apparently not actually proprietary to AT&T, contrary to rumor, though the source code that implements the protocol and is distributed with UNIX is AT&T's trade secret).

Mail is transferred by submitting a mail command over a direct connection by the UUCP remote command execution mechanism [36]. The arguments of that mail command indicate whether the mail is to be delivered locally on that system or resubmitted to another system. In the early days, it was necessary to guess the route to a given host and hope. The only method of acknowledgment was to ask the addressee to reply. Now there is a program (*pathalias*) that can compute reasonable routes from the UUCP map, and there is software that can automatically look up those routes for users.

The UUCP mail network is currently supported in North America mostly by dial-up telephone links. In

Europe there is a closely associated network called EUnet, and in Japan there is JUNET, both of which will be discussed later.

The most common dial-up link speed on the UUCP mail network is 1,200 bps, though there are still a few 300-bps links, and 2,400 bps is becoming more popular. When systems are very close, they are sometimes linked by dedicated lines, often running at 9,600 bps. Some UUCP links are run over local-area networks such as ethernet, sometimes on top of TCP/IP (though more appropriate protocols than UUCP are usually used over such transport media, when UUCP is used its usual point-to-point error correction code is bypassed to take advantage of the reliability of the underlying network and to improve bandwidth). Some such links even exist on long-haul packet-switched networks.

The widespread use of more sophisticated mail relay programs (such as sendmail and MMDf) has increased reliability. Still, there are many hosts with none of these new facilities, and the sheer size of the network makes it unwieldy.

The UUCP mail network has traditionally used source routing with a syntax like `host!hostb!hostc!host!user`. The UUCP map and *pathalias* have made this bearable, but it is still a nuisance. An effort is underway to alleviate the routing problems by implementing naming in the style of ARPA Internet domains. This might also allow integration of the UUCP name space into the ARPA Internet domain name space. In fact there is now an ATT.COM domain in which most hosts are only on UUCP or CSNET. Most UUCP hosts are not yet in any Internet domain, however. This domain effort is also handled by the UUCP Project and appears to be proceeding at a methodical but persistent pace [35, 37].

The hardware used in the UUCP mail network ranges from small personal computers through workstations to minicomputers, mainframes, and supercomputers. The network extends throughout most of North America and parts of Asia (Korea and Israel). Including hosts on the related networks JUNET (in Japan) and EUnet (in Europe), there are at least 7,000 hosts on the network; possibly 10,000 or more. (EUnet and JUNET hosts are listed in the UUCP maps.) The UUCP Project addresses are

`uucp-query@cbatt.ATT.COM`
`cbatt!uucp-query`
`uucp-query@cbatt.UUCP`

Much information about UUCP is published in USENET newsgroups.

USENET

USENET began in 1980 as a medium of communication between users of two machines, one at the University of North Carolina, the other at Duke University. It has since grown exponentially to its current size of more than 2000 machines. In the process the software has been rewritten several times, and the transport mechanisms now used to support it include not only the original UUCP links, but also X.25, ACSNET, and others.

USENET combines the idea of mailing lists as long used on the ARPANET with bulletin-board service such as has existed for many years on TOPS and other systems, adding a freedom of subject matter that could never exist on the ARPANET, and reaching a more varied constituency. While chaotic and inane ramblings abound, the network is quite popular.

The USENET news network is a distributed computer conferencing system [23] bearing some similarities to commercial conferencing systems like CompuServe, though USENET is much more distributed. Users pursue both technical and social ends on USENET. Exchanges are submitted to groups on various topics, ranging from garden astronomy.

The name "USENET" comes from The USENET Association, The Professional and Technical User's Group. The name UNIX is a pun on Multics [65], which is the name of a major predecessor operating system. (The pun indicates that, in areas where Multics tries to do many things, UNIX tries to do one thing well.) USENET has no central administration, though there are newsgroups to which introductory and other information about the network is posted monthly. USENET is currently defined as the set of hosts receiving the newsgroup `news.announce` (but see the EUnet and JUNET sections). There are about a dozen hosts that constitute the backbone of the network, keeping transit times low by doing frequent transfers among themselves and with other hosts that they feed. Since these hosts bear much of the burden of the network, their administrators tend to take a strong interest in the state of the network. Most newsgroups can be posted to by anyone on the network. For others, it is necessary to mail a submission to a moderator, who decides whether to post it. Most moderators just filter out redundant articles, though some make decisions on other grounds. These newsgroup moderators form another group interested in the state of the network. Newsgroups are created or deleted according to decisions made after discussion in the newsgroup `news.groups`.

Each host pays its own telephone bills. The back-

bone hosts have higher bills than most other hosts due to their long-distance links among themselves. The unit of communication is the news article. Each article is sent by a flooding routing algorithm to all nodes on the network [34]. The transport layer is UUCP for most links, although many others are used, including ethernet, berknet, and long-haul packet-switched networks; sometimes UUCP is run on top of the others, and sometimes UUCP is not used at all.

The many problems with USENET (e.g., reader overload, old software, slow propagation speed, and high and unevenly carried costs of transmission) have raised the possibility of using the experience gained in USENET to design a new network to replace it. The new network might also involve at least a partial replacement for the UUCP mail network.

One unusual mechanism that has been proposed to support the new network is Stargate [89, 90]. Commercial television broadcasting techniques use unused bandwidth in the vertical blanking interval between picture frames. Some broadcasters are currently using this part of the signal to transmit text services. Since many cable-television channels are distributed via geosynchronous satellites, a single input to a satellite uplink facility can reach all of North America on an appropriate satellite and through a satellite uplink company interested in carrying USENET-like articles to be broadcast by cable on a well-known cable-television channel.

Prototypes of hardware and software to encode the articles and other hardware to decode them from a cable-television signal have been built and tested in the field for more than a year; a new, reasonably priced model of the decoder may be available soon.

The facility would allow most compatible systems to use the footprint (area of coverage) of the satellite to obtain access to the appropriate cable-television channel to obtain decoding equipment and hook into the cable network at a very reasonable cost. Articles can be submitted for transmission by UUCP links to the satellite uplink facility. Most of the technical problems of Stargate seem to have been solved.

More than 90 percent of all USENET articles reach more than 90 percent of all hosts on the network within three days. Though there have been some famous bugs caused by loss of articles, that particular problem has become rare.

Every USENET host has a name. That host name and the name of the poster are used to identify the source of an article. Though those hosts that are on both the UUCP mail and USENET news networks

usually have the same name on both networks, mail addresses have no meaning on USENET: Mail related to USENET articles is usually sent via UUCP mail; it cannot be sent over USENET, by definition. Though the two networks have always been closely related, there are many more hosts on UUCP than on USENET. In Australia the two networks do not even intersect except at one host.

There are different distributions of newsgroups on USENET. Some go everywhere, whereas others are limited to a particular continent, nation, state or province, city, organization, or even machine, though the more local distributions are not really part of USENET proper. The European network EUnet carries some USENET newsgroups and has another set of its own [79]. JUNET in Japan is similar to EUnet in this regard.

There are about 2000 USENET hosts in the United States, Canada, Australia, and probably in other countries. The hosts on EUnet, SDN, and JUNET communicate with USENET hosts: The total number of news hosts including ones on those three networks is probably at least 2500. The UUCP map includes USENET map information as annotations. A list of legitimate netwide newsgroups is posted to several newsgroups monthly. Volunteers keep statistics on the use of the various newsgroups (all 250 of them) and on frequency of posting by persons and hosts. These are posted to `news.newslists` once a month, as is the list of newsgroups. Important announcements are posted to two moderated newsgroups, `news.announce` and

`news.announce.newusers`, which are intended to reach all users (the current moderator is Mark Horton, `cbsgd!mark`). An address for information on the network is `seismo!usenet-request`.

EUnet

The European UNIX network (EUnet) started at the April 1982 European UNIX systems Users' Group (EUUG) meeting in Paris, and originally connected machines in the Netherlands, Denmark, Sweden, and the United Kingdom. It began as an extension or application of the software and protocols used in USENET and UUCP in North America, and most hosts, as on those networks, run UNIX. There have always been and still are some marked differences, however. Mail and news are much more closely tied together in EUnet: The backbone hosts and administrators are the same for both, and a single name is used for the combined mail and news network. The administration of the network is much more organized than for UUCP and USENET in North America, and there has always been a much stronger rela-

tionship between EUnet and EUUG than there has ever been between USENET or UUCP and either USENIX or /usr/group (the two organizations in North America most similar to EUUG). Many of the "soapbox" discussion groups are not carried on EUnet (due to their high costs in transatlantic and European traffic), and there are many newsgroups that are only distributed within Europe. EUnet also differs from USENET, UUCP, and ACSNET in almost every other regard, especially funding, as will be seen. Yet the purpose of the network, to provide its users with modern communication facilities, particularly electronic mail and news, which are capable of reaching the users of as many networks as possible, is in line with the other networks.

There is one EUnet backbone host in each European member country. Each such backbone host organizes communications within its country, often by maintaining direct connections to all other hosts in the country. The backbone hosts also communicate among themselves across international boundaries. The whole set of backbone hosts is the backbone of the network. There is a central host to which all backbone hosts have connections and that carries all the intercontinental news and most of the mail traffic; this host has always been *mcvax* in Amsterdam.

The administrators of the backbone hosts hold meetings (usually at EUUG meetings) where they determine concerted strategies and tactics. Currently they are moving toward implementing ARPA Internet RFC886 domains in EUnet and expect to have completed doing so by mid 1987. The top-level domains being chosen are the ISO-3166 two-letter country codes, following RFC920. When a single legal or political entity is needed to speak for the network, EUUG does so. There is some financial support from EUUG, but most funds are provided by the owners of the individual hosts. For example, the cost of the news connections with North America is shared proportionately among the EUnet backbone hosts according to the number of news hosts each feeds. They in turn share these costs equally among all the hosts in their country. Thus no host bears a disproportionate burden.

Mail is charged to the originating host on a message-by-message and link-by-link basis. Except this is not possible for intercontinental links and links to nonchargeable networks. The originating host usually brings these charges to the attention of the senders of the mail in some manner.

The European Public Data Networks charge per segment (maximum 64 bytes) for use of X.25. There is also a negligible connection time charge and an

initial connection charge, but the per-packet charge accounts for more than 80 percent of the costs. Rates for X.25 connections are lower in Europe than in North America, and usually lower than equivalent North American telephone rates for similar distances and connection times. But, despite uniformity in charging units, X.25 tariffs vary widely within Europe. Mail originating outside of Europe in some cases must be paid for by a EUnet backbone host, particularly when gatewaying to a national non-charging network. This makes bulk mailings rather annoying to the gateway administrators.

English is the lingua franca of the network, though many other European languages are also used. The basic application protocols are news and mail as in USENET and UUCP, plus remote login where X.25 links are used. EUnet originally used UUCP over dial-up telephone links like those of USENET and UUCP in North America. This arrangement did not last more than about six months, as it quickly became evident that X.25 links were more practical (faster and cheaper) in Europe for long-distance links. Most EUnet links between backbones outside Europe are now UUCP (without much of the usual error checking) over X.25. Ordinary UUCP telephone dial-up links are still the most common for local links and to leaf nodes.

EUnet has recently concluded an arrangement to redistribute the ACSNET software within Europe, and it appears likely that many links between backbone hosts will soon use ACSNET rather than UUCP because of such obvious advantages as multiple links and full duplex connections.

The SLIP software that allows TCP/IP over serial lines may be used over some interbackbone links in place of X.25 because of its advantage in speed (4,800 bps is the practical limit for X.25 links), especially over some of the faster digital telephone links now becoming available (up to 50,000 bytes at about \$0.25 a minute). EUnet will probably eventually adopt X.400 for mail, using whatever transport mechanisms, such as TP4 or other ISO transport protocols, are appropriate at that time.

Most dial-up telephone links are still 1,200 bps, but 2,400 bps is becoming more common. X.25 links are mostly effective about 4,800 bps, though the nominal rate most commonly used is 9,600 bps. Ninety percent of all mail and news traffic arrives within one day. Reliability is quite high.

EUnet has already almost completely eliminated the old-style UUCP host!hostb!host!user syntax in favor of user@host. Routing is managed by the backbone hosts, each of which knows the organization within its own country and which hosts are in

which country. Routing information is automatically exchanged daily between the backbone hosts and between Europe and the United States and Japan. If mail is sent from a nonbackbone host that lacks a direct link to the destination host, it is forwarded to the national backbone host, which relays it.

ARPA Internet RFC886 domain naming syntax is currently being implemented on EUnet. Each country will register as a top-level country domain with the Internet (e.g., NL for the Netherlands). There will also be subdomains from the beginning. Thus *mcvax* might become *cwi.ac.nl*, and an address for a user of that host there might be *user@cwi.ac.nl*. This simplifies routing further, since each backbone host need then only know the hosts within its own country domain and a path to the backbone host for each other country domain: There is no need to know anything about the internal structure of other country domains. In fact, routing is delegated further since the internal structure of a given subdomain may be known by a host other than the national backbone host, which need then only know an appropriate nameserver host for the subdomain. EUnet has about 900 hosts and extends throughout Western Europe, as shown in Table VI. There are connections to EARN, JANET, DFN, and other networks within Europe, plus intercontinental connections to Japan and Korea, to CSNET, UUCP, and INTERNET in North America, and to ACSNET and INTERNET in Australia. New hosts must register with the national backbone host administrator, and inquiries within Europe should be addressed to the EUnet secretariat.

TABLE VI. EUnet Connected Hosts as of July 1, 1986

Austria	7
Belgium	17
Denmark	36
Finland	47
France	94
Germany	90
Greece	1
Ireland	6
Italy	28
Netherlands	129
Norway	7
Sweden	108
Switzerland	50
United Kingdom	276
Total	896

Throughput for June 1986: 550 Mbytes. Total throughput for 1986 of one backbone host: 500 Mbytes; average throughput: 10 Mbytes per month. Split news/mail is about 50/50. Costs: average U.S. \$150 for data transport.

Source: courtesy of Teus Hagen and Piet Beertema.

same place. Elsewhere, direct inquiries to EUUG:

EUUG secretariat
Owles Hall
Buntingford
Herts SG9 9PL
England
mcvax!euug

SDN

The System Development Network (SDN) [10, 11] was started in 1982 in the Republic of Korea with one node at Seoul National University and another at the Korea Institute of Electronics Technology. The major development issues in SDN during the initial period from 1982 to 1984 were setting up the environment for computer communications and adding new nodes. SDN is now the backbone that interconnects local-area networks at major sites. Most of the nodes are connected via the TCP/IP protocol suite through leased lines, and advanced research is being carried out in network softwares, international standards, and distributed systems. The intent of the network is to provide a facility for computer communications and resource sharing, and a test environment for research and development communities in Korea.

Technical and administrative support has been provided by Korea Advanced Institute of Science and Technology (KAIST) since 1983. The Network Management Center, located at KAIST, handles information dissemination and also maintains international and domestic contacts for administrative matters. Managerial decisions are made by the Overseeing Committee, which consists of representatives from each site. The Electronics and Telecommunications Research Institute and the Data Communications Company of Korea, along with KAIST, are the major participants in the management and development of SDN.

Each site is charged the cost for its connections, whether through leased lines, X.25 PDN, or dial ups for domestic communications. International communications costs are charged proportionally. Expenses for protocol development and management and international communication are covered by national research grants, public corporations, and internal funding from several institutes.

The standard protocol architecture in SDN is based on the U.S. DoD IP, and UUCP is also supported over all hosts. Some UUCP links run on top of TCP/IP and X.25. Virtual terminal, file transfer, mail (in both Korean and English), remote command execution, net news, and name services are sup-

ported. The SDN research community is currently working on migration to ISO protocols starting from the network and transport layer. The low-level network is built largely out of leased phone lines and the domestic X.25 network, in addition to LANs. The international connections are based on X.25 and X.28/X.29 PAD. Most links run at 2,400 bps, and some links run at 9,600 bps. Reliability is high.

The top domain name, SDN, has been used for years in SDN. Second-level domains are usually hosts. Routing is decided at the issuing host. In addition to the domain style name, UUCP-like *host/user* names can be used with a *pathalias* database.

Internet domain naming as proposed in RFC882 and RFC883 is being implemented. The naming structure consolidated as of June 1986 is in conformance to RFC920. The top-level domain is KR for Korea, with the following second-level domains below it:

RES for research community
EDU for educational institutes
COM for companies
GOV for governmental organizations
ORG for general organization
<network-names> for other nationwide networks

Third-level domains are the organization names.
Fourth-level domains usually apply to hosts.

UUCP gateways to North America were set up in 1983 over X.25 and dial-up lines. A CSNET link using PMDF over X.25 was installed in 1984. The CSNET connection will soon be replaced by TCP/IP over the X.25 link.

Within the Republic of Korea, about 100 computers at 22 different organizations were connected to SDN as of June 1986. By the end of 1987, 200-300 computers at 30 different organizations will be connected. SDN administrators are now working on Pacnet [9], a cooperative network over UUCP connections that will interconnect universities and companies in the Pacific region. SDN currently has connections to Japan (kddl) and Indonesia (indovax) directly, and to Australia (munnari) and Singapore (tataelxsi) through *seismo*.

SDN liaison is *Kilnam Chon*:

UUCP: *seismo!kaist!chon*
ARPA Internet: *chon%kaist.csnet@relay.cs.net*

The mailing address is

SDN Management Center
Department of Computer Science
Korea Advanced Institute of Science and Technology
P.O. Box 150 Chongryang

Seoul 131

Republic of Korea

UUCP: *seismo!kaist!nmc*

ARPA Internet: *nmc%kaist.csnet@relay.cs.net*

JUNET

The earliest large computer network in Japan appears to have been N-1 [40], which is an interuniversity network that started late in 1970. It currently connects about 20 national universities using DD (Digital Data eXchange, Packet-switching network) which is the Japanese Public Data Network. (This is also a circuit switching network, like Canada's Infoswitch.) N-1 provides only remote login and mail service.

JUNET originally linked three universities starting in October 1984. It was connected to Europe in January 1985 by a link between *kddl* and *mcvar*. Connections to other continents have since been established. It is currently the major nationwide noncommercial computer network in Japan (DDX-P might be better). JUNET incorporates both news (like USENET) and mail (like UUCP) in a single network organization (like EUnet). Its purpose is to promote information exchange among Japanese researchers and with researchers outside Japan [61].

The network is organized by a group of students at the Tokyo Institute of Technology, whose leader is Jun Murai. The administrators of the major (backbone) hosts on the network also help administer the network and hold monthly meetings for that purpose.

There is evidently a tradition in Japan of employees asking their employers for permission before publishing anything publicly; this might explain the small number of Japanese news postings seen outside Japan. Also, there is a distribution (file) of newsgroups that can be seen only inside Japan. For instance, *fi.kanji* (for kanji handling) and *fi.micro.mac* are very active. Each host's connection costs are paid by its institution.

UUCP is the common protocol, with X.25, telephone dial ups, and ethernet below. Some UUCP links are carried over TCP/IP. The ISO protocols are not yet used, and there appear to be no immediate plans to use them, although implementations exist within several corporations. Mail (per RFC822) and news (as in USENET) are supported. There are three standard human-language notations supported:¹

English in 8-bit ASCII:

We had a very good time at Mark's home last night.

¹ Thanks to Hide Takuda and Jun Murai for these examples: JIS is Japan Industrial Standard, and C6228 defines Kanji characters of 2 bytes each of 7 bits.

Japanese in roman characters in 8-bit ASCII:

Watashi tachiwa sakuban Mark no iede tanoshii
hitotoki wo sugoshita.

Japanese in a 16-bit JIS C6228 encoding of Kanji
characters:

私達は、昨晚マークの家で楽しい一時を過ごした。

The mail and news interface software has been modified to support Kanji characters, which also require special display hardware.

Most JUNET links are 2,400 bps, and reliability is high. There is a domain system, and its top-level domain is JUNET. Second-level domains are usually named for organizations such as universities, and third-level domains are usually hosts. This scheme is similar to RFC733 [14], the predecessor of the current ARPA Internet domain system. Routing is done by tables on gateways, which are manually updated once a month. Modifications have been made to sendmail to support routing by domains.

JUNET has 46 participating organizations and 100 hosts. There are now several links to Europe, North America, Australia, and Korea. A gateway to BITNET (BITNET in Japan) is planned, and there is a gateway to CSNET. The contact address is

Y. Murai
Computer Center
Institute of Technology
Yama Meguroku
152 Japan

NETWORKS

is a metanetwork, but has already been described in the section on research networks.

The history of NSFnet properly starts with JUNET and continues with CSNET. In 1984, the National Science Foundation (NSF) established the Office of Advanced Scientific Computing (OASC). The office started one program to develop supercomputer centers, and another to develop network links to them. The planned network was called NSFnet [42]. This network will eventually provide services other than just supercomputer access, and will serve an extensive user community. Agreement was reached between NSF and DARPA in October 1985 to provide access to NSFnet for ARPANET users and vice versa.

The objective of NSFnet is to provide the general

academic community with the kind of networking resources that CSNET now provides to computer science researchers [3]. It extends the ideas of resource sharing, which motivated the ARPANET, and of collaboration among researchers, which grew out of the development of the ARPANET [19]. NSFnet will also concentrate more on customer support and information than has been customary in the ARPA Internet.

The NSFnet project is administered by the OASC in coordination with the administrations of existing networks. The NSF will provide initial funding, although there is potential for the network to eventually become self-supporting.

Diverse network protocols will be used to support the TCP/IP protocol suite; the network may eventually migrate to the ISO protocols. Initially, NSFnet will use the existing long-haul (ARPANET) and campus networks of the ARPA Internet plus state networks and a new supercomputer backbone network. NSFnet can be viewed as an expansion of the ARPA Internet: The ARPANET alone may gain about 15 new PSNs. BITNET is considering migrating from RSCS to TCP/IP, and CSNET is investigating a new network, Cypress, based on TCP/IP 9,600-bps leased lines. Both of these may be used as part of NSFnet. The TCP/IP protocols are also being implemented under the Cray CTSS operating system.

In addition to services such as already exist in CSNET and the ARPA Internet, it is hoped that NSFnet will have a sophisticated nameservice extending not only to hosts, but also to users and resources. A program might be given a networkwide location-independent name, such as *Useful-Program*, which would be mapped by the nameservice into *uprog@somedomain*. Users would not have to be concerned with the location of the program.

Long-haul links are expected to typically be T1 microwave links. Reliability will be high due to the use of TCP/IP. Naming, addressing, and routing will be as in the ARPA Internet. Administrators of CSNET, BITNET, and UUCP have recently agreed to adopt the ARPA Internet domain name system (though conformance is voluntary per host in the latter two networks). This decision was motivated by purely practical considerations, but will facilitate the implementation of NSFnet.

The network will reach academic campuses and research organizations throughout the United States and already has 65 initial participant organizations. For further information, contact

Program Director for Networking
Office of Advanced Scientific Computing
National Science Foundation
Washington, DC 20050

National Research Internet

The National Research Internet (NRI) is being planned by a network working group of several White House committees. The relevant committees are collectively known as the Federal Coordinating Committee on Science, Engineering, and Technology (FCCSET), and their purpose is to coordinate activities in the various federal agencies. One of them is the FCCSET Committee on Very High Performance Computing (VHPC), which is aimed at the supercomputer activities.

NRI will interconnect the networks owned by the various federal agencies in order to promote resource sharing and collaboration among researchers. This differs from NSFnet's purpose in that it attempts to satisfy most if not all needs for scientific networking, whereas NSFnet is more directed toward providing access to supercomputer centers. NSFnet will be one of the networks connected by NRI. The NRI project is administered by the FCCSET Network Working Group, and the network will be funded by various federal agencies. Accounting methods are still under study.

The TCP/IP protocol suite will be used for NRI. The main focus of development will be on gateway standards in order to facilitate interconnection of the networks. The links used will presumably be fast, and the TCP/IP protocols will promote reliability.

Access restrictions among the various agency networks may become an issue, and prior arrangement with the administration of a network may be necessary before a given resource on one network can be used by a user of a different network.

RARE

The EEC's ESPRIT project is sponsoring plans for a metanetwork called RARE (Réseaux Associées pour la Recherche Européenne). RARE is an attempt to unify and standardize the European national networks. Funding may come from both the EEC and from the governments of participating countries. The ISO protocols will be used, though ESPRIT also sponsors projects like ROSE that may convert existing networks from other protocols gradually. Components of some international networks, such as EUnet or EARN, may be funded and incorporated. Some constituents of this metanetwork already exist in national or regional networks like DFN (West Germany) or NORDUNET (Scandinavia). NORDUNET uses not only ISO protocols, but also some Coloured Book ones.

Pacnet

Pacnet is a logical grouping of Pacific hosts and organizations [9]. It is currently in the early planning

stages, though there are some existing UUCP links between the various national networks in the Pacific region, including East Asia, Australia, and the United States.

AUSEAnet

AUSEAnet is a metanetwork for a joint microelectronics (VLSI) project among ASEAN (Association of South East Asian) countries and Australia. The ASEAN countries include Thailand, Indonesia, Malaysia, Singapore, Brunei, and the Philippines. The project started in July 1986 and is expected to be operational by November 1986. The goal is to prepare electronic submission of VLSI designs to the fabrication plant in Australia and to exchange information about microelectronics techniques. Funds are provided by the Australian government and will be augmented by the other participant countries. AUSEAnet will use UUCP and ACSNET over the international X.25 networks. Outside of Australia, it will use mostly UUCP. Most of the links will be 1,200 bps. Reliability should be high.

Country domain and node addressing mode will be decided in August 1986. Indonesia will act as the ASEAN regional center of AUSEAnet and will connect to a designated node in Australia through an international packet-switching line. Each participating ASEAN country will have its own network connected to an international gateway that will poll the Indonesian hub machine. The national center in Indonesia will be at NETLAB, the Network Laboratory, the part of the Inter-University Center for Computer Science at the University of Indonesia in Jakarta. Four institutional organizations were connected to NETLAB as of July 1986. At least nine ASEAN institutional nodes will participate in AUSEAnet. The person is Joseph F. P. Luhukay, who can be reached at

Computer Science Center
University of Indonesia
P.O. Box 3442
Jakarta 10002
Indonesia

BULLETIN BOARDS AND NETWORKS

Bulletin-board systems have become ubiquitous, and many people today think of bulletin-board systems when they hear the word "network." The two are in fact quite distinct. A typical bulletin-board system is a personal computer supporting a database of messages. Users can read or submit messages by dialing up one or more modems over the public switched telephone network. The telephone network is the only network involved, since all the messages are on one machine.

Bulletin-board systems are computer conferencing systems. Unlike mail systems, where communication is one-to-one or one-to-many, communication on computer conferencing systems is many-to-many [32]. People who use personal-computer bulletin-board systems have a tendency to call anything larger a network. An example of this is the WELL (Whole Earth 'Electronic Link) [6], which is a VAX-11/750 with two Fujitsu Eagle disk drives running UNIX (4.2BSD) in the office of the Whole Earth Review in Sausalito, California. Despite the difference between the WELL's 800 megabytes of disk space and the typical home computer's few megabytes, the WELL is still a bulletin-board system. People dial it up and use it like they would CompuServe or THE SOURCE. The WELL can do more than the average BBS system, however. There is a local magazine that is published on-line, as well as parts of other magazines and books [7]. In addition, the WELL is on USENET and UUCP and serves as one of the few public access entry points to those networks. While USENET users are employees of organizations that have machines connected to USENET, anyone can dial up and submit a valid credit-card number (in principle) be a WELL user. Access is not limited, however, since the administrators of the WELL are aware of the nonprofit nature of USENET. There are several other public access UNIX systems that provide USENET access, such as the Soup in New Jersey. Few of them support public adventures, however.

Commercial Networks

Commercial networks sell services to outside users. Many are in effect common carriers like a telephone system. Administration is always centralized, though execution may be delegated. Fees are usually charged to individual persons or organizations on the basis of connect or CPU time used. CompuServe, THE SOURCE, and other such systems are not really networks. They consist of a large number of computers closely coupled into a large distributed system and are accessed just like home personal-computer bulletin-board systems, except that users get bills. In a commercial service, more traffic is good, because it brings in more revenue, but, on an anarchic network like USENET, more traffic just means more congestion. Users of either kind of system have to deal with information overload, though the availability of money and administration in the centralized service allows the development of sophisticated filtering mechanisms to limit that problem. But centralization produces its own problems: During popular hours, most of the dial-up ports and PDNs—public

data networks—that are used to reach them become saturated, and the mainframes themselves become loaded.

PDNs are real networks, or at least communication subnets, as they are implemented up to the network layer and enable users to reach computers. There are many PDNs; every European country, for instance, has a national government-run PDN. The government agency involved is usually the same one that runs the paper post and telephone services. The French acronym PTT (Poste, Téléphone, et Télégraphe) is the standard abbreviation for this kind of agency.

Conferencing Networks

A basic technical difference that distinguishes conferencing systems from mail systems is that conferencing systems need multicast or broadcast transport mechanisms to support them, whereas mail systems can get along with unicast mechanisms. This and other technical issues related to the extent and method of distribution of the message database are not highly visible to the users of any of the BBS or news systems.

Local bulletin boards have existed for many years on TOPS-20 (and other) machines on the ARPA Internet, with one message database per system. Those systems pass the messages in the databases around to each other so that the database on each system approximates those on the others. The transfer methods are an ad hoc combination of file transfer of the message databases and mail of individual messages. There are many established ARPANET mailing lists, many of which are stored on arrival at TOPS-20 systems as BBOARDS, but these are still not true distributed systems.

USENET is similar in concept to the TOPS-20 BBOARD systems. Each USENET host machine contains one copy of the news article database. The key difference is that, when a USENET article is posted locally, it is sent to every neighboring machine, which in turn sends it to all of its neighbors, and so forth. To keep this flooding algorithm from producing loops, each article contains a record of the path it has traversed, which is updated at each machine the article passes through; no machine will send an article to a machine that is already in the path. The article also contains a unique ID that the hosts use to recognize and discard duplicates. Many people have taken to having their personal computers dial up CompuServe in the middle of the night (when telephone rates are low) and download many articles for later perusal. This is a step in the USENET direction. Most computer conferencing systems organize messages into conferences according to subject mat-

ter. Many have each conference (SIG, newsgroup, etc.) overseen by a person who may be called an editor, a moderator, or a Sysop (the term used on most commercial and personal systems). This person filters out duplicate submissions, and may in some cases reject objectional submissions or remove them after they are posted. Reasons for rejection vary widely according to the network, conference, and people involved. Sometimes actual editing is done. The role of the moderator—as perceived by the moderators themselves, as well as network administrators, submitters, and readers—can vary widely. For example, accusations of censorship are unfortunately common, though few moderators believe they are justified.

SOCIAL AND LEGAL ISSUES

Networks have effects on their users that go beyond the issue of immediate practical utility [32]. The primary effect is increased human interaction. If information overload can be avoided, increased interaction can lead to better technical productivity through the exchange of ideas and references. It can also lead to unanticipated social interaction among specialized or diverse groups of people. In addition to social effects, network interaction has potential legal implications for users and administrators.

Networked Communities

Certain newsgroups, mailing lists, bulletin boards, and SIGs have reliable followings that form social groups. These range from groups interacting strictly in pursuit of technical goals to others interacting for the sake of interaction, to still others for whom the networked interaction is an aspect of or leads to outside interaction.

Technical Groups.

UNIX-WIZARDS@BRL.ARPA: This ARPA Internet mailing list² dates back to around 1977 on the ARPANET and is currently gatewayed bidirectionally and automatically with the USENET newsgroup `comp.unix.wizards`. It is possible that most working UNIX software developers and system administrators read this list up to a few years ago, but many have since canceled their subscriptions because of how long it takes to sort through the much larger volume of submissions. There have been several attempts to reduce the traffic and to keep it more technical. The `comp.unix.help` newsgroup, which is gatewayed with the `INFO-UNIX@BRL.ARPA` mailing list, was created to provide novices with access to

² The convention for subscribing to ARPA Internet mailing lists is to send mail to `list-REQUEST@domain`, not `list@domain`. For example: If you want to get on the HUMAN-NETS list, mail a request to `HUMAN-NETS-REQUEST@RED.RUTGERS.EDU`. Only actual submissions should go to `HUMAN-NETS@RED.RUTGERS.EDU`.

knowledgeable people, and so to keep elementary questions out of UNIX-WIZARDS. This new newsgroup is popular, at least. There is also a moderated newsgroup, `comp.unix`. This has little traffic, apparently because people do not want to have to justify the value of their submissions. UNIX-WIZARDS still has a recognizable group of technical contributors and readers who use it in their work. Many of them can also be found attending USENET conferences for the same reasons. Many of those who no longer follow UNIX-WIZARDS use other newsgroups or mailing lists or private mail for the same purpose.

AILIST@SRI-AI.ARPA or `comp.ai`: This is a general technical discussion list or newsgroup for artificial intelligence (AI) researchers. It is moderated and digested. The volume is high, and topics range from press treatment of AI to esoteric points of logic implementation details. Submitters range from the most eminent practitioners to novices, with the moderator selecting more for the former. It is unclear that this list accurately reflects the working community, but it certainly has its own following.

TCP-IP@SRI-NIC.ARPA: This is an ARPANET mailing list that deals with the TCP/IP protocol suite. It is used both for dissemination of information to people not familiar with the protocols, and for working technical discussions among their mentors, most of whom appear to follow the list. There are other similar lists on more specific working topics.

news.group: This is a USENET newsgroup for discussing the creation and deletion of newsgroups. It has on occasion been one of the highest volume newsgroups on the network. There are other newsgroups that are also about USENET itself.

INFO-NETS%MIT-OZ@MC.LCS.MIT.EDU: This is a mailing list about networks. Typical postings might deal with requests for paths to specific hosts on certain networks or requests for position statements by people involved with NSFnet. Some information in this article was obtained from responses to requests on this list.

HUMAN-NETS@RED.RUTGERS.EDU: HUMAN-NETS is perhaps the prototypical technical list about social issues. It is a forum for discussions of the social effects of computers and specifically of computer networks. A discussion in this list led to the writing of an earlier version of this article.

There are technical mailing lists for such things as workstations, local-area networks, and many different lists for many different manufacturers' computers. Not all technical lists or newsgroups are

computer related. There are newsgroups on astronomy and biology, for instance. However, researchers in other fields do not use newsgroups in their fields for actual work as much as researchers in computer fields do, probably because researchers in other fields are less familiar with unusual uses of computers.

Social Groups.

CompuServe SF SIG: CompuServe has a very popular though not very old special interest group (SIG) in science fiction, moderated in part by Diane Duane, a popular science-fiction and fantasy writer. The instigators had great difficulty convincing network management that a conference on this topic would be viable, but it has turned out to be one of the fastest growing SIGs. Similar groups exist on some other systems.

There are many newsgroups or mailing lists that are not only for social purposes. A famous example among aficionados was a mailing list started by a student who had lost his girlfriend and wanted to commiserate with all his friends, most of whom he knew through the various networks. This list used considerable portions of the bandwidth of several networks over many months and led to a number of outages in several parts of North America where the participants met each other directly. This list was eventually sanctioned by the administrators of any network.

Effects

One of the most obvious effects of networks is their tendency to induce users to "flame," that is, to produce many words on an uninteresting topic, or in an angry or ridiculous manner; "raving" is almost a synonym for flaming. The usual explanation for why computer networks tend to aggravate flaming is that the flamer is isolated from the readers and has no immediate feedback to inhibit such behavior. There are typographic conventions that have developed on the various networks to get around the limitations of expressing nuances in ASCII characters. One of the more universal is that UPPERCASE means shouting (much to the chagrin of those with keyboards that only have uppercase). Some "surround phrases with asterisks" to indicate emphasis, while others space the characters out. People will mark sarcasm or irony. Facial expressions often get spelled out ("grin"). There are many ways to indicate the start of a flame, such as *FLAME ON!*. On USENET there are shorter ways to indicate lack of malicious intent: for instance, :-) (the image of a smiling face).

Legal Issues

The specific liabilities that arise when computers communicate with other computers over public airways or through the telephone system can be difficult to recognize. The extent of a system's liability is a function of its specific features and of the extent to which the content of its messages is controlled. There are legal precedents covering the liabilities of more traditional communication media such as newspapers, radio and television broadcasting, and cable television. Though computer networks do not fit neatly into either the broadcaster or common-carrier category, these are the two classifications that seem likely to provide legal precedents that will apply to computer communications [78]—the alternative is to define a special classification for computer communications. The classification is important because common carriers are not held to as strict a standard as broadcasters. Some of the liabilities faced by network administrators have to do with defamatory material, obscenity, content of transmission, and faulty transmission. Individual users might also be liable for defamatory material and obscenity, as well as for copyright infringement and invasion of privacy.

To control what is posted on a network, it is necessary to control access to that network. Most existing networks are not strong on security. The safest policy in using networks is to assume that any network can be broken, that any transmission can be recorded, and that most can be forged. (There was a famous hoax on April Fools' Day in 1984 when kremvax!kgbvax!chernenko joined USENET.) Encryption techniques for providing a rather high degree of security exist, but few people are willing to pay the price in CPU time, and so few networks make use of these techniques.

HISTORY

The first packet-switching network was implemented at the National Physical Laboratories in the United Kingdom. It was quickly followed by the ARPANET in 1969 [15]. There were soon related projects such as CYCLADES in France [72], EIN (European Informatics Network) in Europe [2, 21], and the Coloured Book efforts in the United Kingdom. Of these, the networks based on the ARPANET and Coloured Book protocols have survived, by mutation and evolution of their protocols. These original efforts were all undertaken as research and were largely government supported.

As the idea of networking caught on, companies such as Xerox, DEC, and IBM started to develop their own networking technologies, usually starting with local-area networks. Long-haul networks came

to be used not only for communication among directly connected hosts, but also to tie LANs into internets.

ARPANET technology had been used by BBN to build the commercial network TELENET (later sold to GTE) by 1976, and commercial X.25-based networks followed. In Europe the PTTs controlled (and still control) the PDNs in each country (one per country) and have universally settled on X.25 as their network layer protocol. The PTTs favor circuit switching rather than packet switching, so most of the CCITT protocols such as X.25 and X.400 are oriented toward virtual circuits.

Computer conferencing systems started in 1976 and later found commercial viability in centralized services such as EIES and Delphi as well as somewhat more distributed systems like CompuServe. Personal computers are often used as free bulletin boards.

Meanwhile, another networking technology, based on dial-up telephone links instead of dedicated connections, was being developed. Two of the earliest products of this technology were ACSNET and UUCP, both of which survive in modified forms. The dial-up networks produced the most distributed of the conferencing systems: USENET.

CSNET started as an attempt to bring the collaborative advantages of the ARPANET to researchers beyond the ARPANET community by using dial-up mechanisms similar to those of UUCP. MAILNET grew out of a similar effort. BITNET made IBM's internal mainframe networking technology available to the academic community and even spread to some non-IBM hosts. It has also spread outside the United States as NETNORTH and EARN.

Internets required new protocol suites, such as Xerox NS, the U.S. DoD's TCP/IP, and the ISO protocols. The spread of NS has, some say, been stifled by the secrecy of its originating company. The TCP/IP protocols are by far the most widely implemented of these three due to the accessibility of their specifications, their long history of practical use, and the backing of the U.S. government. Some of the ISO protocols found implementation in 1983 on CDNnet, the first EARN network, and spread rapidly in Europe the following year. Other implementations, particularly of X.400 (actually a CCITT protocol), have followed, especially in Europe. Most of the ISO protocols are either adoptions of CCITT protocols or are, like them, oriented toward virtual circuits.

Hosts on early networks were usually either mainframes or minicomputers. A few networks, such as BITNET, continue this tradition. Internets usually


have many workstations on their LAN components so the average size of their hosts is smaller. Personal computers are sometimes connected to internets in the ARPA Internet, and some are part of the dial-up networks. Users of IBM PCs have found a network of their own in FidoNet. At the other end of the spectrum, at least one network, MFENET, was developed primarily for access to supercomputers.

Computer networks have spread to larger and smaller machines, different lower layer technologies, different protocols, and many nations. Though the diversity continues to increase, most noncommercial networks are connected with each other at least for the purposes of mail exchange and thus already constitute a worldwide metanetwork, first predicted years ago and called Worldnet (see Figure 2).

BIBLIOGRAPHIC NOTES

A good introduction to actual functioning networks is [50]. The standard tutorial introduction to the theory of computer networks is still [85]. The reader can gain a useful amount of context by reading its first chapter, and there is an overview paper by the same author [86]. Either of these should be supplemented by recent publications on the various protocols and protocol suites, for which see the specific references. A perspective on historical and recent developments in protocol design and implementation can be found in [66]. Good discussion concerning networks and protocols can be found in [38], and there are some useful comparative network papers in [13]. Much of the information in this article was obtained over the networks it describes either by mail, file transfer, or news. Some of the references are to articles in well-known digests or newsgroups.

Acknowledgments. The number of people who have supplied information for or otherwise helped with this article is much too large for their names to be listed here; the authors thank all of them. However, we would like to especially acknowledge a few whose aid has extended far beyond their own networks and the call of duty, including Rick Adams, Janet Asteroff, Piet Beertema, Kilnam Chon, Pete Collinson, John R. Covert, Franklin Davis, Peter J. Denning, Robert Elz, Erik Fair, Michel Gien, Mark Horton, Jin H. Hur, Christian Huitema, Daniel Karrenberg, Steve Kille, Joshua Knight, John Larson, Barry Leiner, Jun Murai, Henry Nussbacher, Craig Partridge, Hans Strack-Zimmermann, Gligor Tashkovich, Valerie Thomas, Hide Tokuda, and doubtless others whom we have inadvertently neglected to mention. Many network administrators



Solid lines indicate a continuing operational network, and dashed lines indicate a transfer of technology to a newer network. Italics indicate internets.

FIGURE 2. Time Lines for the Development of Notable Computer Networks

have supplied information, but information included about any network may not necessarily reflect the views of its administrators, users, or others associated with it. Inclusion or omission of a section on any network does not imply anything about the importance, quality, size, or any other property of that network: The networks included are simply those that the authors noticed. Any intentional inaccuracies or other faults in this article are the sole responsibility of the authors.

REFERENCES

1. Allman, E. Interview with Peter Honeyman. In *UNIX Review*, vol. 4, no. 1, M. Compton, Ed. Freeman, San Francisco, Calif., Jan. 1986, p. 64.
2. Barber, D.L.A. A European informatics network: Achievement and prospects. In *IEEE '76* (Toronto, Canada, Aug.). IEEE, New York, 1976, pp. 44-50.
3. BBN Laboratories. 4th anniversary issue. *CSNET News 8* (Summer 1985). CSNET CIC, BBN Laboratories, Cambridge, Mass.
4. Blumann, W., Cadwallader, R., Diedi, A., Lovelock, J., Power, D., Pozzana, S., and Saury, C. Implementation of OSI protocols in the ESPRIT information exchange system. In *ESPRIT '85: Status Report of Continuing Work*. Elsevier North-Holland, New York, 1986, pp. 1387-1398.
5. Boggs, D.R., et al. Pup: An internetwork architecture. *IEEE Trans. Commun. COM-28*, 4 (1980), 612-624.
6. Brand, S. Gate Five Road. *Whole Earth Rev.* 47 (July 1985), 103-104.
7. Brand, S. Gate Five Road. *Whole Earth Rev.* 48 (Fall 1985), 139.
8. Cerf, V.G., and Cain, E. The DoD internet architecture model. *Comput. Networks* 7 (1983), 307-316.
9. Chon, K. National and regional computer networks for academic and research communities in the Pacific region. In *Proceedings of PCCS* (Seoul, Korea, Oct.). 1985, pp. 560-566.
10. Chon, K., et al. System development network. In *Proceedings of TENCON* (Singapore, Apr.). 1984, pp. 133-135.
11. Chon, K., et al. SDN: A computer network for Korean research community. In *Proceedings of PCCS* (Seoul, Korea, Oct.). 1985, pp. 567-570.
12. Comer, D. The computer science research network CSNET: A history and status report. *Commun. ACM* 26, 10 (Oct. 1983), 747-753.
13. Compton, M., Ed. *Unix Review*, vol. 4, no. 1. Freeman, San Francisco, Calif., Jan. 1986.
14. Crocker, D.H., Vital, J.J., Pogran, K.T., and Henderson, D.A., Jr. Standard for the format of ARPA network text messages: RFC733. In *ARPANET Working Group Requests for Comments*. SRI International, Menlo Park, Calif., Nov. 21, 1977.
15. DARPA. *A History of the ARPANET: The First Decade*. Bolt, Beranek and Newman, Cambridge, Mass., Apr. 1983. (Defense Tech. Info. Ctr. AD A1 15440.)
16. Davis, F. "Worldnet" and X.400 in Switzerland. IAM, Universitaet Bern, Bern, Switzerland, Apr. 1986. ünibe@sys.ife.ethz.chunetæ
17. Denneit, S., Feinler, E.J., and Perillo, F. *ARPANET Information Brochure*. DDN Network Information Center, SRI International, Room E[291, 333 Ravenswood Ave., Menlo Park, CA 94025, Dec. 1985.
18. Denning, P.J. Internal structure of networks. *Am. Sci.* (Mar.-Apr. 1985).
19. Denning, P.J. Supernetworks. *Am. Sci.* (May-June 1985), 225-227.
20. Denning, P.J., Hearn, A.C., and Kern, C.W. History and overview of CSNET. In *Proceedings of the ACM SIGCOMM Symposium* (Mar.). ACM, New York, 1983, pp. 138-145.
21. Deparis, M., Duenki, A., Gien, M., Louis, J., LeMoli, G., and Weaking, K. The implementation of an end-to-end protocol by EIN centres: A survey and comparison. In *Proceedings of IEEE '76* (Toronto, Canada, Aug.). IEEE, New York, 1976.
22. Dick-Lauder, P., Kummerfeld, R.J., and Elz, R. ACSNET—The Australian alternative to UUCP. In *USENIX Association Conference Proceedings* (Salt Lake City, Utah, June 12-15). USENIX Association, El Cerrito, Calif., 1984, pp. 11-17.
23. Emerson, S.L. USENET: A bulletin board for UNIX users. *Byte* (Oct. 1983), 219.
24. Fuchs, I.H. BITNET: Because it's time. *Perspect. in Comput.* 3, 1 (Mar. 1983).
25. Gilmore, P., and Neufeld, G. A strategy for a national electronic messaging system for research in Canada. In *Proceedings of the Canadian Information Processing Society National Meeting* (Montreal, Canada, June), 1985.
26. Ginsberg, K. Getting from here to there. In *Unix Review*, vol. 4, no. 1, M. Compton, Ed. Freeman, San Francisco, Calif., Jan. 1986, p. 45. ("How UNIX communications facilities compare with the routing mechanisms of TYMNET, DNA, SNA, and ARPANET—and why UNIX users should care.")
27. Green, J.L., and King, J.H. Behind the scenes during a comet encounter. *EOS* 67, 105 (Mar. 1986).
28. Green, J.L., and Peters, D., Eds. Introduction to the space physics analysis network (SPAN). TM-86499. NASA Ames Research Center, Moffett Field, Calif., Apr. 1985.
29. Green, J.L., and Zwickl, R.O. Data system users working group meeting report. *EOS* (meeting report) 67, 8 (1986).
30. Green, J.L., Peters, D.J., Heijden, N., and Lopez-Swafford, B. Management of the space physics analysis network. NSSDC Publication July 1986.
31. Hendricks, E.C., and Hartman, T.C. Evolution of a virtual machine subsystem. *IBM Syst. J.* 18, 1 (1979), 111.
32. Hiltz, S.R., and Turoff, M. Structuring computer-mediated communication systems to avoid information overload. *Commun. ACM* (July 1985), 580-589.
33. Hinden, R., and Sheltzer, A. The DARPA internet gateway. In *ARPANET Working Group Requests for Comments*, no. 823. SRI International, Menlo Park, Calif., Sept. 1982.
34. Horton, M.R. Standard for interchange of USENET messages. RFC850. In *ARPANET Working Group Requests for Comments*. SRI International, Menlo Park, Calif., June 1983.
35. Horton, M.R. What is a domain? In *USENIX Association Conference Proceedings* (Salt Lake City, Utah, June 12-15). USENIX Association, El Cerrito, Calif., 1984, pp. 368-372.
36. Horton, M.R. UUCP mail interchange format standard. In *ARPANET Working Group Requests for Comments*, no. 976. SRI International, Menlo Park, Calif., Feb. 1986.
37. Horton, M.R., Summers-Horton, K., and Kercheval, B. Proposal for UUCP/Usenet registry host. In *USENIX Association Conference Proceedings* (Salt Lake City, Utah, June 12-15). USENIX Association, El Cerrito, Calif., 1984, p. 373.
38. IEEE. *Proceedings of the 25th IEEE Computer Society International Conference (COMPCON)*. Society Press, Los Angeles, Calif., Sept. 1981.
39. International Organization for Standardization. ISO open system interconnection—Basic reference model. ISO/TC 97/SC, vol. 1, 719. International Organization for Standardization, Aug. 1981. Ordering information: American National Standards Institute, 1430 Broadway, New York, NY 10018. Phone: 1-212-642-4900; elsewhere are equivalent national ordering offices.
40. Ishida, H. Current status of the N-1 inter-university network with access to supercomputers in Japan. In *Proceedings of PCCS* (Seoul, Korea, Oct.). 1985.
41. Jacobsen, O., and Postel, J. Protocol document order information. RFC980. In *ARPANET Working Group Requests for Comments*, no. 980. SRI International, Menlo Park, Calif., Mar. 1986.
42. Jennings, D.M., Landweber, L.H., Fuchs, I.H., Farber, D.J., and Adrien, W.R. Computer networking for scientists. *Science* 231 (Feb. 28, 1986), 943-950.
43. Kawaguchi, K., Sato, K., Sample, R., Demco, J., and Hilpert, B. Interconnecting two X.400 message systems. In *Proceedings of the 2nd International Symposium on Computer Message Systems* (Sept.). FIP, Washington, D.C., 1985, pp. 15-26. (Also in *Computer Message Systems*—85, pp. 17-28.)
44. Kiessig, R. *UUCP Network Directory*, vol. 2, no. 2. Spring 1986. P.O. Box 50174, Palo Alto, CA 94303.
45. Kille, S.E. Mapping between X.400 and RFC822; RFC987. In *ARPANET Working Group Requests for Comments*, no. 987. SRI International, Menlo Park, Calif., June 1986.
46. Kluger, L., and Shoch, J. Names, addresses, and routes. In *Unix Review*, vol. 4, no. 1, M. Compton, Ed. Freeman, San Francisco, Calif., Jan. 1986, p. 30.
47. Kummerfeld, R.J. *ACSnet: Current Status and Future Development*. Korea, 1985.
48. Kummerfeld, R.J., and Dick-Lauder, P.R. The Sydney Unix network. *Aust. Comput. J.* 13, 2 (May 1981), 52-57.
49. Kummerfeld, R.J., and Dick-Lauder, P.R. Domain addressing in SUN III. In *Proceedings of EUUG* (Paris). 1985.
50. Landweber, L.H., Jennings, D.M., and Fuchs, I. Research computer networks and their interconnection. *IEEE Commun. Mag.* 24, 6 (June 1986), 5-17.
51. Leiner, B.M., Cole, R., Postel, J., and Mills, D. The DARPA internet

- protocol suite. *IEEE Commun. Mag.* (Mar. 1985). (Also in 1985 DDN Protocol Handbook, vol. 2.)
12. McQuillan, J.M., and Walden, D.C. The ARPA network design decisions. *Comput. Networks* 1 (1977), 243-289.
 13. MIL-STD. *Military Standards for DoD Internet Protocols*. Naval Publications and Forms Center, Code 3015, 5801 Tabor Ave., Philadelphia, PA 19120. The ARPA Internet protocols are defined by the set of military standards IP (MIL-STD-1777), TCP (MIL-STD-1778), FTP (MIL-STD-1780), SMTP (MIL-STD-1781), and TELNET (MIL-STD-1782). See also ARPANET Working Group Requests for Comments and 1985 DDN Protocol Handbook.
 14. Mills, D.L. Exterior gateway protocol formal specification; RFC904. In ARPANET Working Group Requests for Comments, no. 904. SRI International, Menlo Park, Calif., Apr. 1984.
 15. Mockapetris, P. Domain names—Concepts and facilities; RFC882. In ARPANET Working Group Requests for Comments, no. 882. SRI International, Menlo Park, Calif., Nov. 1983.
 16. Mockapetris, P. Domain names—Implementation and specification; RFC883. In ARPANET Working Group Requests for Comments, no. 883. SRI International, Menlo Park, Calif., Nov. 1983.
 17. Mockapetris, P. The domain name system. In *Proceedings of the IFIP 65 Working Conference on Computer Message Services* (Nottingham, England, May). IFIP, Washington, D.C., 1984. (Also as ISI/RS-84-133, June 1984.)
 18. Mockapetris, P. Domain system changes and observations; RFC973. In ARPANET Working Group Requests for Comments, no. 973. SRI International, Menlo Park, Calif., Jan. 1986.
 19. Mockapetris, P., Postel, J., and Kirton, P. Name server design for distributed systems. In *Proceedings of the 7th International Conference in Computer Communication* (Sydney, Australia, Oct.), 1984. (Also as ISI/RS-84-132, June 1984.)
 20. Mogul, J., and Postel, J. Internet standard subnetting procedure; RFC950. In ARPANET Working Group Requests for Comments, no. 950. SRI International, Menlo Park, Calif., Aug. 1985.
 21. Mural, J., and Asami, T. A network for research and development communities in Japan—JUNET. In *Proceedings of PCCS* (Kaist, Korea, 1985), pp. 579-588.
 22. Nield, G., Demco, J., Hilpert, B., and Sample, R. EAN: An X.400 message system. In *Proceedings of the 2nd International Symposium on Computer Message Systems* (Sept.), IFIP, Washington, D.C., 1985.
 23. (Also in *Computer Message Systems*—85, pp. 3-15.)
 24. D.A. Uucp implementation description. In *UNIX Programmer's Manual*, 7th ed., vol. 2. Holt, Rinehart and Winston, New York, 1987.
 25. D.A., and Lesk, M.E. A dial-up network of UNIX systems. In *Programmer's Manual*, 7th ed., vol. 2. Holt, Rinehart and Winston, New York, 1983.
 26. FEL. *The Multics System: An Examination of Its Structure*. Addison-Wesley, Reading, Mass., 1975.
 27. F.M.A. *The Elements of Networking Style*. Prentice-Hall, Englewood Cliffs, N.J., 1985.
 28. Bee, C. Mail routing and the domain system; RFC974. In ARPANET Working Group Requests for Comments, no. 974. SRI International, Menlo Park, Calif., Jan. 1986.
 29. Bee, C. Report from the Internet NIC on domains. CSNET-MAIL Dig. 2, 2 (Feb. 19, 1986). CSNET CIC, BBN Laboratories, Cambridge, Mass.
 30. Bee, C. Mail routing using domain names: An informal tour. In *Proceedings of the 1986 Summer USENIX Conference* (Atlanta, Ga., June 11-14), USENIX Association, El Cerrito, Calif., 1986, pp. 366-377.
 31. Bee, C. Domain name system implementation schedule—Revised; RFC921. In ARPANET Working Group Requests for Comments, no. 921. SRI International, Menlo Park, Calif., Oct. 1984.
 32. Bee, C., and Reynolds, J. Domain requirements; RFC920. In ARPANET Working Group Requests for Comments, no. 920. SRI International, Menlo Park, Calif., Oct. 1984.
 33. Buzin, L., Ed. *The CYCLADES computer network—Towards layered network architectures*. In *Monograph Series of the ICCG*, vol. 2. North-Holland, New York, 1982, p. 387.
 34. Quarterman, J.S., Silberschatz, A., and Peterson, J.L. 4.2BSD and BSD as examples of the UNIX system. *ACM Comput. Surv.* 17, 4 (Dec. 1985), 379-418.
 35. Richie, D.M., and Thompson, K. The UNIX time-sharing system. *Comput. Syst. Tech.* 1, 57, 6, Part 2 (July-Aug. 1978), 1905-1929. (The original version (*Commun. ACM* 17, 7 (July 1974), 365-375.) described Version 6, whereas this one describes Version 7.)
 36. Anderson, T., Ho, S., Heijden, N., Jabs, E., and Green, J.L. Near-realtime data transmission during the ICE-comet Giacobini-Zinner encounter. *ESA Bull.* 45, 21 (1986).
 37. Seamonson, L.J., and Rosen, E.C. "STUB" Exterior gateway protocol; RFC888. In ARPANET Working Group Requests for Comments, no. 888. SRI International, Menlo Park, Calif., Jan. 1984.
 38. Shoch, J.F. Internetwork naming, addressing, and routing. In *Proceedings of the 17th IEEE Computer Society International Conference (COMPCON)* (Sept.), IEEE, New York, 1978, pp. 430-437.
 39. Shulman, G.H. Legal research on USENET liability issues. *login: The USENIX Assoc. Newsl.* 9, 5 (Dec. 1984), 11-17. (USENIX Association, El Cerrito, Calif.)
 40. Simonsen, K.J. Re: Re: Network differences. *net.unix*, no. 15@duku.UUCP, USENET, Dec. 18, 1985.
 41. SPAG. *Guide to the Use of Standards*. 1985.
 42. Spratt, E.B. Networking developments in the U.K. academic community. In *International Conference on Information Network and Data Communication (INDC-86)*, IFIP TC.6 (Ronnedy, Sweden, May). IFIP, Washington, D.C., 1986.
 43. SRI International. 1985 DDN Protocol Handbook. DDN Network Information Center, SRI International, Room E1291, 333 Ravenswood Ave., Menlo Park, CA 94025, 1985. Most relevant papers from both ARPANET Working Group Requests for Comments and Military Standards for DoD Internet Protocols are collected here.
 44. SRI International. ARPANET Working Group Requests for Comments. DDN Network Information Center, SRI International, Room E1291, 333 Ravenswood Ave., Menlo Park, CA 94025. This series of technical notes includes the specifications for the ARPA Internet protocols IP (RFC791), ICMP (RFC792), TCP (RFC793), UDP (RFC768), FTP (RFC959), SMTP (RFC821), and TELNET (RFC854), plus related papers. All the protocols are indexed in Assigned Numbers (RFC960) and Official ARPA-Internet Protocols (RFC961).
 45. Summers-Horton, K., and Horton, M. Status of the USENIX UUCP project. In *USENIX Association Conference Proceedings* (Dallas, Tex., Jan. 23-25). USENIX Association, El Cerrito, Calif., 1985, p. 183.
 46. Tanenbaum, A.S. *Computer Networks*. Prentice-Hall Software Series, Prentice-Hall, Englewood Cliffs, N.J., 1981.
 47. Tanenbaum, A.S. Network protocols. *ACM Comput. Surv.* 13, 4 (Dec. 1981), 453-489.
 48. Uhlig, R., Ed. *Computer Message Systems*—85. Elsevier North-Holland, New York, 1986.
 49. VNET-BITNET Gateway General User's Guide. Available on BITNET via INFO@BITNIC, Mar. 22, 1985.
 50. Weinstein, L. Project Stargate. In *USENIX Association Conference Proceedings* (Portland, Oreg., June 11-14), USENIX Association, El Cerrito, Calif., 1985, pp. 79-80.
 51. Weinstein, L. Project Stargate. In *EUUG Conference Proceedings* (Florence, Italy, Apr.). European Unix Systems Users' Group, Buntingford, England, 1985.
 52. Xerox Corporation. The Ethernet, a local area network: Data link layer and physical layer specification. X3T51/80-50, Xerox Corporation, Stamford, Conn., Oct. 1980.
 53. Xerox Corporation. An internetwork architecture. XSI 028112, Xerox Corporation, Stamford, Conn., Dec. 1981.

CR Categories and Subject Descriptors: C.2.0 [Computer-Communication Networks]: General—data communications; K.2 [History of Computing]: systems, networks
General Terms: Management, Performance, Reliability, Standardization

Additional Key Words and Phrases: Addressing, domains, gateways, networks, protocols, Worldnet

Author's Present Address: John S. Quarterman, P.O. Box 14621, Austin, TX 78751; Josiah C. Hoskins, MCC, 9430 Research Blvd., Kaleido Bldg. 2, Austin, TX 78759-6509.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.